



JOHANNES STEGMAIER

New Methods to Improve Large-Scale Microscopy Image Analysis with Prior Knowledge and Uncertainty

Johannes Stegmaier

New Methods to Improve Large-Scale Microscopy Image
Analysis with Prior Knowledge and Uncertainty

New Methods to Improve Large-Scale Microscopy Image Analysis with Prior Knowledge and Uncertainty

by
Johannes Stegmaier

Dissertation, Karlsruher Institut für Technologie (KIT)
Fakultät für Maschinenbau

Tag der mündlichen Prüfung: 3. Juni 2016

Hauptreferent: apl. Prof. Dr.-Ing. Ralf Mikut

Korreferenten: Prof. Dr. Uwe Strähle, Prof. Dr. Jan G. Korvink

Impressum



Karlsruher Institut für Technologie (KIT)
KIT Scientific Publishing
Straße am Forum 2
D-76131 Karlsruhe

KIT Scientific Publishing is a registered trademark of Karlsruhe
Institute of Technology. Reprint using the book cover is not allowed.

www.ksp.kit.edu



*This document – excluding the cover, pictures and graphs – is licensed
under the Creative Commons Attribution-Share Alike 3.0 DE License
(CC BY-SA 3.0 DE): <http://creativecommons.org/licenses/by-sa/3.0/de/>*



*The cover page is licensed under the Creative Commons
Attribution-No Derivatives 3.0 DE License (CC BY-ND 3.0 DE):
<http://creativecommons.org/licenses/by-nd/3.0/de/>*

Print on Demand 2017 – Gedruckt auf FSC-zertifiziertem Papier

ISBN 978-3-7315-0590-7

DOI 10.5445/KSP/1000060221

Zusammenfassung

Jüngste Entwicklungen im Bereich der mehrdimensionalen Mikroskopie bieten ein großes Potential für die Beantwortung vielerlei Fragestellungen in wissenschaftlichen Bereichen. Beispielsweise bieten neue Verfahren wie die zeitaufgelöste 3D Konfokal- und Lichtscheibenmikroskopie oder die Transmissions-elektronenmikroskopie weitreichende Möglichkeiten im Bereich der Biologie, die von der ganzheitlichen Analyse embryonaler Entwicklung über die Betrachtung subzellulärer Prozesse bis hin zur Rekonstruktion von Verschaltungen im Nervensystem von Modellorganismen reichen. Die dabei routinemäßig anfallenden Datenmengen im Terabyte-Bereich können allerdings nur unzureichend manuell ausgewertet werden und eine wichtige Komponente für die erfolgreiche Auswertung solcher bildbasierter Experimente ist daher eine größtmögliche Anzahl von Analyseschritten durch Bildanalyseverfahren zu automatisieren. Bestehende Verfahren für die automatische Bildauswertung sind hierbei jedoch meist nicht unmittelbar auf die großen Datenmengen anwendbar und die Analysen müssen daher auf kleine Auszüge der Daten beschränkt werden, falls die enormen Anforderungen an Rechenleistung und Verarbeitungszeit nicht gewährleistet werden können. Zudem wird vorhandenes *a priori* Wissen oft nur unzureichend in automatische Verfahren eingebettet und somit ein bedeutender Teil an Zusatzinformationen vernachlässigt, der zu einer verbesserten Ergebnisqualität beitragen kann.

Die Hauptbeiträge der vorliegenden Arbeit sind ein neues Konzept zur Abschätzung und Weiterleitung von Unsicherheiten in Bildverarbeitungsketten sowie die Entwicklung neuer Segmentierungsverfahren, die für eine effiziente Analyse von 3D Mikroskopbildern im Terabyte-Bereich eingesetzt werden können. Basierend auf unscharfen Mengen (engl. fuzzy sets) wurde zur Verfügung stehendes Vorwissen systematisch in eine mathematische Repräsentation überführt, die anschließend für effiziente Datenselektion, eine Unsicherheitsabschätzung von automatisch extrahierten Daten sowie für die gezielte Verbesserung von Bildverarbeitungsoperatoren eingesetzt werden konnte. Um

den Bedarf an effizienten Bildverarbeitungsalgorithmen zu reduzieren wurden drei neue Segmentierungsalgorithmen entwickelt, die sich für eine Extraktion von sphärischen, linienförmigen und lokal planaren Objekten eignen. Die neuen Segmentierungsmethoden wurden dabei gezielt für den Einsatz in der automatisierten Analyse von großen 3D Bilddatensätzen optimiert, insbesondere durch die systematische Ausnutzung von vorhandenem *a priori* Wissen während der Algorithmenentwicklung und durch die Beschränkung auf rechen- und speichereffiziente Teilkomponenten in der Implementierung. Anhand einer exemplarischen Bildverarbeitungskette wurde veranschaulicht, wie Unsicherheiten in bestehende Operatoren integriert werden und zur Verbesserung der Ergebnisqualität beitragen können. Um die Funktionalität der vorgestellten Verfahren zu validieren, wurden erweiterte oder zum Teil neu erstellte, simulierte Benchmarkdatensätze verwendet, die eine Vielzahl möglicher Einsatzszenarien systematisch abdecken. Die effizienten Implementierungen werden zum einen innerhalb der vorliegenden Arbeit vorgestellt und zum anderen als plattformunabhängige Open-Source Software zur allgemeinen Verfügung bereitgestellt. Eine Reihe von Problemen im Bereich der Entwicklungsbiologie wurde mittels der theoretisch eingeführten Verfahren erfolgreich ausgewertet. So wurden die Methoden beispielsweise für eine automatisierte, quantitative Analyse der Auswirkungen von bekannten und unbekanntem Chemikalien auf die neuronale Entwicklung im Rückenmark von Zebrafischarmblingen, für die Detektion, Segmentierung und zeitliche Verfolgung von fluoreszenzmarkierten Zellkernen in Zebrafischarmblingsembryos sowie zur quantitativen Charakterisierung von Zellmorphologieänderungen in zeitaufgelösten 3D Mikroskopbildern von Fruchtfliegen-, Zebrafischarmblings- und Mausembryos eingesetzt.

Abstract

Recent developments in the area of multidimensional imaging techniques provide powerful ways to examine various kinds of scientific questions. For instance, in biological applications, time-resolved 3D light-sheet microscopy and serial section electron microscopy provide unprecedented possibilities ranging from *in toto* analyses of embryonic development down to investigations of sub-cellular processes or reconstructions of the nervous system. The routinely produced datasets in the terabyte-range, however, can hardly be analyzed manually. Thus, the extensive use of image analysis-based automation is an essential key to the success of the performed imaging experiments. Existing algorithms for such analysis tasks are mostly not directly applicable to these large-scale datasets and either have to be confined to small excerpts of the data or require an immense amount of computation capacities and execution time. Moreover, available prior knowledge that could be exploited for advanced analyses is often not sufficiently considered by automatic processing pipelines.

The major contributions of the present thesis are a new concept for the estimation and propagation of uncertainty involved in image analysis operators and the development of new segmentation algorithms that are suitable for terabyte-scale analyses of 3D+t microscopy images. Based on fuzzy set theory, available *a priori* knowledge was transformed into a mathematical representation and extensively used to enhance the performance of processing operators by data filtering, uncertainty propagation and explicit exploitation of information uncertainty for result improvements. To target the need for efficient image analysis operators, three new segmentation algorithms were specifically developed to detect a generalized geometric class of objects, namely, spherical objects, line-like objects and locally plane-like objects. The developed pipelines were specifically tuned to be applicable to large-scale analyses, *i.e.*, only fast and memory efficient processing operators were used in the implementation. Using an exemplary pipeline, it is demonstrated how a combination of both the

fast algorithms and the proposed uncertainty framework could be used to further enhance the overall quality of the considered processing operators. All developed methods were thoroughly validated on existing and newly developed simulated benchmarks, to be able to quantitatively assess their applicability to different imaging conditions. In addition, the efficient implementations of all developed algorithms are presented and were made accessible to the community as platform independent open-source software tools. The new methods were successfully applied to multiple large-scale analyses of fluorescence microscopy images in the field of developmental biology. In particular, the proposed pipelines were used to quantify the impact of both known and unknown chemical substances on the neuronal development in the spinal cord of zebrafish in 2D images. Furthermore, the developed methods were applied to time-resolved 3D images to detect, segment and track fluorescently labeled cellular nuclei of entire zebrafish embryos and to quantitatively characterize cell morphology dynamics using fluorescently labeled cellular membranes in 3D+t microscopy images of fruit fly, zebrafish and mouse embryos.

Acknowledgements

In the first place I want to thank Prof. Dr.-Ing. habil. Georg Bretthauer for the great opportunity to spend my time as a PhD student at the Institute for Applied Computer Science (IAI) at the Karlsruhe Institute of Technology (KIT) and for his supervision of the thesis. Special thanks to my direct supervisor apl. Prof. Dr.-Ing. Ralf Mikut for the guidance, encouragement, constructive discussions and the continuous support throughout the entire time at the IAI. I deeply appreciate having had the chance to work freely and responsibly on a highly captivating project. I want to thank Prof. Dr. Uwe Strähle and Prof. Dr. Jan G. Korvink for reviewing the thesis and Prof. Dr. Barbara Deml for heading the examination board. Thanks to all colleagues, Bachelor's students and trainees at the IAI, especially, to Rüdiger Alshut, Thomas Anritter, Andreas Bartschat, Dr. Christian Bauer, Wolfgang Doneit, Eduard Hübner, Arif ul Maula Khan, Jorge Angel Gonzalez Ordiano, Nico Peter, Willis Pinaud, Dr. Markus Reischl, Benjamin Schott, Manuel Traub, Michele Rene Tuga and Simon Waczowicz for the pleasant, cooperative and exciting working atmosphere in the group for biosignal analysis. Many thanks to all the collaboration partners from the Institute of Toxicology and Genetics (ITG) and the Institute for Applied Physics (APH), especially, for igniting my interest in developmental biology. In particular I want to thank Dr. Thomas Dickmeis, Dr. Andrei Kobitski, Prof. Dr. G. Ulrich Nienhaus, Dr. Jens C. Otte, Dr. Sepand Rastegar, Dr. Maryam Shahid, Prof. Dr. Uwe Strähle, Dr. Masanari Takamiya, Dr. Benjamin Weger, Dr. Meltem Weger and Dr. Lixin Yang for all the exciting projects I was allowed to contribute to during the time being at KIT. I really appreciate all the infrastructural support such as data storage management and cluster computing received by Serguei Bourov, Dr. Ariel García, Volker Hartmann, Dr. Rainer Stotzka, Jos van Wezel and all the persons behind the scenes that made the required large-scale analyses technically feasible. For the support, guidance and exciting advanced training courses I want to thank my thesis advisory committee (apl. Prof. Dr.-Ing. Ralf Mikut, Dr. Markus Reischl, Dr. Ute

Schepers and Prof. Dr. Uwe Strähle) and the BioInterfaces International Graduate School (BIF-IGS). Moreover, I'm grateful for the unique opportunity to do an inspiring internship at Howard Hughes Medical Institute's Janelia Farm Research Campus (JFRC). For all their support and supervision during the internship I want to thank Dr. Philipp J. Keller, Dr. Fernando Amat and the whole Keller Lab. For financial support during my time as a PhD student at KIT and during the internship at JFRC, I want to thank the Helmholtz Association in the program BioInterfaces, the Howard Hughes Medical Institute (HHMI) and the Karlsruhe House of Young Scientists (KHYS). For their contributions to the implementation and improvements of XPIWIT I want to thank Dr. Fernando Amat (fast GPU implementations), Andreas Bartschat (XML functionality, filter implementations and pipeline workflow), Eduard Hübner (graphical user interface) and apl. Prof. Dr.-Ing. Ralf Mikut (Gait-CAD interface). Additionally, I want to thank Rüdiger Alshut, Andreas Bartschat, apl. Prof. Dr.-Ing. Ralf Mikut, Dr. Sebastian Pfeiffer, Dr. Markus Reischl and Karl Wöll for their contributions to the IMVid and Tracking toolboxes. Last but not least, I want to cordially thank my girlfriend, family and friends for their continuous support and understanding when I was buried in thoughts once in a while.

Contents

Zusammenfassung	i
Abstract	iii
Acknowledgements	v
1 Introduction	1
1.1 Theoretical Background and Related Work	3
1.1.1 Image Acquisition	3
1.1.2 Image Analysis	5
1.1.3 Benchmarking	15
1.1.4 Prior Knowledge and Uncertainty	16
1.1.5 Available Software Solutions	17
1.2 Open Questions	18
1.3 Objectives and Thesis Outline	20
2 Uncertainty Estimation and Propagation in Image Analysis Pipelines 23	
2.1 The Image Analysis Pipeline Concept	24
2.2 Identification of Suitable Prior Knowledge	25
2.3 Prior Knowledge-based Uncertainty Quantification	26
2.3.1 Quantifying Prior Knowledge using Fuzzy Set Membership Functions	28
2.3.2 Combination of Fuzzy Set Membership Functions	32
2.4 Uncertainty Propagation in Image Analysis Pipelines	33

2.4.1	Uncertainty-based Object Rejection	34
2.4.2	Extended Information Propagation to Compensate Operator Flaws	35
2.4.3	Resolve Ambiguities using Propagated Uncertainty	36
2.4.4	Improved Processing Pipeline by Uncertainty Propagation	37
3	Efficient Segmentation in Multidimensional Image Data	39
3.1	Seed Point Detection	41
3.1.1	Validation Benchmark	42
3.1.2	Seed Detection using a Laplacian-of-Gaussian Maximum Projection	43
3.1.3	Seed Detection using Thresholding and Euclidean Distance Maps	46
3.1.4	Validation	47
3.1.5	Discussion	50
3.2	Accurate Extraction and Comparison of Elongated Shapes in 2D Images	51
3.2.1	Validation Benchmark	52
3.2.2	Algorithmic Design	53
3.2.3	Validation	59
3.2.4	Discussion	62
3.3	Efficient Segmentation of Roundish Objects	63
3.3.1	Validation Benchmark	63
3.3.2	Algorithmic Design	64
3.3.3	Validation	68
3.3.4	Discussion	72

3.4	Fast and Accurate Segmentation of Locally Plane-like Structures in 3D Images	74
3.4.1	Validation Benchmark	74
3.4.2	Algorithmic Design	76
3.4.3	Validation	82
3.4.4	Discussion	85
4	Enhancing Algorithms with Uncertainty Treatment	89
4.1	A New Comprehensive Validation Benchmark	90
4.1.1	Simulation of Fluorescently Labeled Objects	92
4.1.2	Generating the Benchmark Images	93
4.1.3	Performance Assessment	98
4.2	Seed Point Detection	101
4.2.1	Improved Detection and Fusion of Redundant 3D Seed Points	101
4.2.2	Extending Seed Detection Algorithms by Uncertainty Handling	102
4.2.3	Validation	104
4.3	Segmentation	109
4.3.1	Extending Segmentation Algorithms with the Uncertainty Framework	110
4.3.2	Uncertainty Guided Segmentation Performance Improvement	112
4.3.3	Validation	114
4.4	Extended Multiview Information Fusion	119
4.4.1	Uncertainty-based Fusion of Extracted Objects	120
4.4.2	Validation	121

4.5	Tracking	125
4.5.1	Resolving Tracking Errors using Propagated Information	126
4.5.2	Validation	127
4.6	Discussion	131
5	New Implementations and Numerical Tools	133
5.1	XPIWIT - XML Pipeline Wizard for the Insight Toolkit	133
5.1.1	XML Pipeline Creation	134
5.1.2	Data Generation	136
5.1.3	Special Filters	136
5.1.4	Gait-CAD Compatibility	137
5.1.5	Graphical User Interface for Rapid Prototyping	137
5.1.6	Implementation Details	138
5.1.7	Comparison to Existing Software Solutions	138
5.2	Extensions of the Open-Source MATLAB Toolbox Gait-CAD	140
5.2.1	The ImVID Extension Package	140
5.2.2	The Tracking Extension Package	141
5.2.3	The SpinalCord Extension Package	141
5.2.4	The Embryo3DT Extension Package	142
5.2.5	The Benchmark Extension Package	143
5.2.6	Semi-Automatic Uncertainty-based Image Analysis	144
6	Automated Quantitative Analysis of Embryonic Development	147
6.1	Automated Quantification of Neuronal Patterns in the Spinal Cord of Zebrafish	147
6.1.1	Dataset Description	148
6.1.2	Automated Image Analysis Framework	149

6.1.3	Results	150
6.1.4	Discussion	152
6.2	Analysis of Embryonic Development in 3D+T Microscopy Images	154
6.2.1	Dataset Description	154
6.2.2	Automated Image Analysis Framework	155
6.2.3	Results	161
6.2.4	Discussion	163
6.3	Automated Segmentation of Fluorescently Labeled Membrane Images	166
6.3.1	Dataset Description	166
6.3.2	Automated Analysis Framework	167
6.3.3	Results	167
6.3.4	Discussion	171
7	Conclusions	173
A	Nomenclature and Symbols	179
B	Infrastructure	189
B.1	Large Scale Data Facility for Data Storage and Processing	189
B.2	Hadoop Streaming to Parallelize XPIWIT	190
C	Performance Assessment	195
C.1	Seed Point Detection	195
C.2	Segmentation	196
C.3	Tracking	202
C.4	Evaluation Platform	203
D	Benchmark Datasets and Parameters	205

Contents

List of Figures	209
List of Tables	213
List of Listings	215
Bibliography	217

1 Introduction

Three-dimensional (3D) imaging techniques have lately revolutionized the possibilities of in-depth analyses in almost any area of natural sciences. Techniques like radar, tomography, microscopy and other image-based techniques allow to probe various unanswered questions in unprecedented detail, *e.g.*, in the fields of material science, physics, biology or medicine [5, 152, 160]. An inherent problem to volumetric imaging, however, is the tremendous amount of data that is produced by these techniques. Especially, experiments that involve a spatio-temporal acquisition of volumetric image data (3D+t) easily produce datasets reaching the terabyte-range. Besides the infrastructural challenges, datasets of this size render manual investigations almost impossible and require an extensive use of automatic image processing and analysis for the successful assessment of experimental results. A multitude of algorithms for automatic image analysis have been presented in the past. However, most of these methods were developed and tested solely on 2D images and their immediate application to 3D is not possible in many cases. In addition to implementation-based hurdles, many mathematically elegant solutions require an immense amount of computing power and computation time to generate results, and are thus often not practically applicable. Furthermore, available prior knowledge is mostly not systematically incorporated into the algorithms, *i.e.*, a significant amount of extra information is wasted instead of being used to improve algorithmic performance.

The aim of this work was to derive a generic and efficient way to enhance image analysis pipelines by an extensive use of prior knowledge and to assess the uncertainty of produced results with respect to predefined, prior knowledge-based criteria. The generally introduced concept was successfully used to filter erroneous detections, to improve flawed results and to generally resolve ambiguities occurring in automatic processing operators, and had only a low impact on the computation time. Moreover, three new segmentation algorithms were developed in order to provide efficient new ways to extract various geometric

shapes such as line-like, spherical or locally plane-like objects from large-scale 3D images. The performance improvements were achieved by maximizing the amount of prior knowledge used for the algorithmic development. Furthermore, the involved processing operators were carefully selected and optimized with respect to both time complexity and memory efficiency, to make large-scale analyses feasible. All methods were validated on extended and newly developed artificial benchmarks, in order to proof their robustness against disturbances and their applicability in a variety of different imaging scenarios. In all scenarios, the algorithms provided excellent results comparable to state-of-the-art algorithms, while lowering the processing times by at least an order of magnitude compared to existing methods. Finally, all methods were used for quantitative analyses of multidimensional fluorescence microscopy images from the field of developmental biology. The presented approaches were successfully used to (1) quantitatively analyze the impact of known and unknown small molecules to neuron populations in the spinal cord of zebrafish embryos, (2) to detect, segment and track fluorescently labeled cell nuclei in terabyte-scale 3D+t image data of zebrafish embryos and (3) to obtain a quantitative characterization of cell shape dynamics in time-resolved 3D microscopy images of fruit fly, zebrafish and mouse embryos.

The next section introduces related work and the context of the thesis, such as fluorescence microscopy, the required image analysis components as well as previous approaches to quantify and use uncertain information in image-based processing pipelines. Furthermore, a summary of existing benchmarks and software solutions is provided. Although most of the developed methodology is generally applicable to problems observed in 3D imaging experiments, the presentation of recent methods and the considered applications mainly focuses on the automatic analysis of 3D fluorescence microscopy images as demanded by the biological application field considered throughout this thesis. An overview of open questions, aims of the thesis as well as a coarse outline of the major attempts how the problems were solved is provided at the end of this chapter.

1.1 Theoretical Background and Related Work

1.1.1 Image Acquisition

In the beginning of image-based experiments the object of interest has to be recorded using the microscopy technique of choice, such as stereo, compound, phase contrast, confocal, light-sheet or electron microscopy to name but a few. Each of the methods has its own benefits and drawbacks and the appropriate technique has to be carefully selected based on the questions to be answered. A great overview of various microscopy approaches can be found in [165]. In this thesis, the main emphasis is put on fluorescence microscopy with a special focus on imaging biological specimens using light-sheet fluorescence microscopy, as all application examples considered in this work stem from the field of fluorescence microscopy.

In principle, fluorescence microscopy represents a special area of light microscopy. Based on illumination at a specific excitation wavelength, emitted fluorescence can be imaged through the detection optics using a special filter that is matched to the emission spectrum of the fluorophore the investigated specimen has been labeled with [159]. The emitted and filtered fluorescent signal can then be captured using detectors like charge-coupled devices (CCD) or complementary metal oxide semiconductor-based detectors (CMOS) and is subsequently digitized for archival, processing and analysis [165]. Besides fluorescent staining methods like Hoechst or DAPI dyes to label nuclear DNA and immunofluorescence approaches using antibodies [58, 171], the use of fluorescent proteins has unveiled enormous possibilities for detailed image-based studies of gene expression and protein targeting in intact cells and whole organisms [231]. Fluorescent proteins like the green or red fluorescent protein (GFP, RFP) can be fused to a gene of interest and are subsequently detectable upon expression of this particular gene in the transgenic target organism using fluorescence microscopy [159].

Approaches such as confocal microscopy and the recently established light-sheet microscopy perform an optical sectioning of the investigated probe to obtain detailed 3D image stacks, possibly of the entire specimen. Moreover, the improved imaging speed of these techniques enables the acquisition of time-series of 3D images, *i.e.*, to obtain 3D videos of dynamic objects [11, 85, 243].

The basic principle of a light-sheet microscope is illustrated in Fig. 1.1. Based on a focused laser light-sheet, sections of a few micrometers are illuminated at a time and captured by orthogonally arranged detection optics [86]. The microscope is adjusted such that only the focused plane is illuminated, to minimize fluorescence in out-of-focus areas, *i.e.*, to effectively reduce light scattering, photobleaching of fluorophores as well as phototoxicity [106]. A 3D image comprised of sequentially acquired planes can be generated by moving the specimen through the light-sheet (Selective Plane Illumination Microscopy, SPIM, [86]) or vice versa by scanning the laser light-sheet through the specimen (Digital Scanned Laser Light Sheet Fluorescence Microscopy, DSLM, [98]). Problems such as light scattering and absorption can be successfully compensated by multiview acquisition schemes that rotate the specimen [86, 96] or use multiple detection paths [112, 230]. Furthermore, a more homogeneous illumination and a reduced light absorption of the imaged plane can be obtained by the use of double-sided illumination [85, 96, 106].




Figure 1.1: Fundamental principle of light-sheet fluorescence microscopy. The specimen is located in a tube-like holder in the center of the microscope and is illuminated using a movable thin sheet of laser light. Illumination is performed orthogonal to the detection objective and is positioned, such that only the focal plane is excited. Three-dimensional images can be acquired with this technique by moving the light-sheet or the sample (reproduced with permission from [85]).

Successfully imaged model organisms using both confocal and light-sheet microscopy are embryos of nematodes (*Caenorhabditis elegans*) [11], fruit flies (*Drosophila melanogaster*) [93, 230], crustaceans (*Parhyale hawaiiensis*) [17], zebrafish

(*Danio rerio*) [93, 96, 106, 136, 143, 158] and mice (*Mus musculus*) [5, 216]. Tab. 1.1 provides an impression of the typical amount of images produced in 3D+t imaging experiments of different model organisms in the field of developmental biology using confocal or light-sheet microscopy. The acquired 3D+t images routinely produced by these imaging systems easily reach the terabyte-range. Since the first commercial system has become available (Zeiss Z1) and detailed building plans for self-assembled microscopes have been published recently [78, 180], the amount of image data produced by light-sheet microscopes will further increase. Most commercially available platforms, however, only provide solutions for archival and viewing of data but offer only limited image analysis capabilities, *i.e.*, there is an urgent need for an elaborate toolbox of automatic image analysis approaches.

Author	Modality	Organism	Duration	Size (Per Image/Embryo)
Bao <i>et al.</i> (2006) [11]	Confocal	Nematode	6.6 h	10.2 MB / 4.0 GB
Tomer <i>et al.</i> (2012) [230]	SiMView	Fruit Fly	20 h	5.1 GB / 11.0 TB
Keller <i>et al.</i> (2008) [96]	DSLIM	Zebrafish	24 h	3.9 GB / 15.3 TB
Kobitski <i>et al.</i> (2015) [106]	DSLIM	Zebrafish	16 h	5.1 GB / 9.0 TB
Amat <i>et al.</i> (2014) [5]	SiMView	Mouse	2 h	332 MB / 8.4 GB

Table 1.1: Exemplary datasets that were used to study the early embryonic development of various model organisms using confocal or light-sheet microscopy. Acquisition durations range from a few hours up to days and produce 3D+t image datasets in the terabyte-range for a single embryo (values taken from [5, 11, 96, 106, 230]).

1.1.2 Image Analysis

Digital Image Representation and the Processing Pipeline Concept

Digital images are usually represented as a 2D, 3D or more generally an N-D matrix, if multiple channels, slices or time points of a scene are captured. Each of the matrix entries directly correlates with the amount of photons that were absorbed by the respective sensors and the analog signal is quantized using an analog-to-digital converter, *i.e.*, a mapping of an analog signal to a finite set of positive integer or floating point values is performed. Common dynamic ranges used for digital images are 8 bit (256 intensity levels), 16 bit (65536 intensity levels) or 32 bit (floating point, usually scaled to the interval $[0, 1]$). For a 16-bit image this corresponds to 0 being the minimum intensity (black) and

65535 being the maximum intensity (white) [34]. Consequently, RGB color images store a separate intensity value for each of the color components red, green and blue. In cases where the physical detector operates with a lower dynamic range (*e.g.*, 12 bit), only a subset of the actually available intensity range provided by the image format is used. On the other hand, if the number of photons hitting a pixel during the exposure time exceeds its full-well capacity, a saturation of the pixel occurs and should be avoided by any means. Of course, the used bit depth directly affects the physical size of the datasets, *i.e.*, converting an 8 bit image to 16 or 32 bit results in a doubling or even quadrupling of the input image and might require compression strategies to tackle the large file sizes. This becomes especially important when dealing with large-scale multi-dimensional image data on an infrastructure with limited data storage capacity, network bandwidth or main memory on the processing nodes. However, some data storage policies require the archival of uncompressed, raw original data, *i.e.*, in these cases it is inevitable to directly preserve the acquired image data as it was produced by the respective microscope and to provide the required infrastructure.

With respect to image filtering, neighborhood relations based on pixel connectivity often have to be considered [88]. Common ways to define the pixel connectivity in 2D images are the 4-neighborhood (pixels that share an edge with the central pixel) and the 8-neighborhood (all pixels surrounding the central pixel). Analogously, the voxel connectivity for 3D images can be defined using a 6-neighborhood (voxels that share a face with the central voxel) or a 26-neighborhood (all voxels surrounding the central pixel). As these definitions fail in border regions of the image, different methods like padding with constant values, the intensity of the closest border pixel or a cyclic mirroring of the image content can be used to allow neighborhood operations in these areas [34].

An important measure to quantify the quality of a digital imaging system is its signal-to-noise ratio. As the name already indicates, this measure connects the signal intensity observed for the investigated objects to the respective back-

ground noise. For the simulated benchmark data considered in this thesis, the following formulation was used to estimate the signal-to-noise ratio:

$$\text{SNR} = \frac{\mu_{fg}}{\sigma_{bg}}, \quad (1.1)$$

with μ_{fg} being the foreground mean intensity and σ_{bg} being the standard deviation of the background signal [209, 236, 246]. Using the available ground truth image as a mask, the bright foreground and low intensity background regions could be optimally separated. An alternative measure to assess the image quality is given by the contrast-to-noise ratio (CNR), which is defined as:

$$\text{CNR} = \frac{|\mu_{fg} - \mu_{bg}|}{\sigma_{bg}}, \quad (1.2)$$

with μ_{fg} being the foreground mean intensity, μ_{bg} being the background mean intensity and σ_{bg} being the standard deviation of the background signal [246]. Although, many variations of these criteria exist, the definition in Eq. (1.1) apparently is the most widely used formulation and was thus used to assess the image quality throughout in this thesis. For background mean intensities close to zero, the two measures yield similar values.

A further relevant aspect to consider is the handling of metadata. For instance, information about the imaged objects, the acquisition procedure, the physical size of the voxels, image resolution and the like, can be used to adjust algorithmic parameters and to convert object properties such as size or volume from image space to physical units. Metadata is often packed into the file header or put into a separate metadata file [3].

After the images and the associated metadata found their way to a digital representation, the processing of images as well as the information extraction is usually formulated as a set of linearly arranged processing operators that modify the image or retrieve information about its content. An overview of exemplary processing steps in a sequential order is shown in Fig. 1.2. Information passed between the processing steps can be images, extracted features or both of them. Of course, some filters may also require the input of multiple preceding processing operators. However, for simplicity, these special cases were omitted in the figure. In the following sections each of the mentioned process-

ing operators is described in more detail and an overview of available methods to accomplish the respective analysis tasks is provided therein.



Figure 1.2: Exemplary image analysis pipeline consisting of image preprocessing, registration and multiview fusion, seed point detection, segmentation and tracking.

Image Preprocessing

Image preprocessing represents an important step in the automated image analysis pipeline of microscopy images. Mostly, the acquired images are flawed by noise, illumination variations, dirt, limited resolution or simply contain unnecessary image regions that need to be discarded. To cope with such flawed image material, the general idea of preprocessing is to enhance and transform the images such that the further processing is facilitated. Using an appropriate preprocessing of the images can have a significant impact on the achieved processing quality of subsequent processing steps such as the seed detection, segmentation and tracking as described in the next sections. One of the most important enhancements is the suppression of noise in the images, *e.g.*, using Gaussian low-pass filtering (Fig. 1.3B), mean filtering, median filtering, quantile-based noise removal or more advanced methods such as anisotropic diffusion filtering or objectness filters that selectively emphasize desired structures based on image derivatives [7, 64, 177]. Besides image noise, illumination inhomogeneities and contrast adaptations can be optimized using techniques like surfaces fitting, histogram equalization, high-pass filtering, morphological filtering or prior knowledge-based feedback techniques [88, 101, 208, 237]. However, if quantitative comparisons of intensity values need to be performed, it has to be ensured that the applied transformations preserve the comparability of different images. To emphasize properties such as edges, *i.e.*, strong local intensity changes, methods like the Sobel filter, the Canny edge detector (Fig. 1.3C), the Laplacian-of-Gaussian (LoG, Fig. 1.3D) or the gradient magnitude can be used [36, 75, 142]. Another frequently used family of efficient preprocessing methods are morphological operations such as erosion, dilation and their combinations like opening, closing and the top-hat filter, which are

perfectly suited to improve imperfect raw images for a specific application [21, 75, 210]. Based on binary structuring elements that define the influence range of the operators, morphological operations can be used for image region separation, filling of holes, structure smoothing, efficient noise reduction, illumination correction and even for gradient approximations [63, 103, 138, 210]. In some cases, only a specific part of the image is important for further consideration. Based on manual selection, intensity distributions, geometrical properties or template matching, the images can be trimmed to the requested region [44, 176].




Figure 1.3: Exemplary results of various image analysis operators. Based on a raw input image of stained nuclei of a zebrafish embryo (A), the individual panels show the results of a Gaussian low-pass filter (B), a Sobel edge detector (C) and a Laplacian-of-Gaussian filter (D). In (E), a binarized version of the input image is shown and the image was used to calculate an Euclidean distance map (F), seed point locations (G) and the uniquely labeled connected components (H). All panels were generated using the open-source software Fiji [200].

Registration and Image Fusion

In many applications, images of different sources have to be spatially aligned to each other, such that the same pieces of information are located at the same places. Popular applications that require image registration are image stitching [33], movement estimation [202], the combination of biomedical images from different modalities [81], the registration of multiple views obtained from 3D

microscopy [183, 185, 230] or the alignment of tissue sections in connectomics projects [47]. In the scope of this thesis only rigid registrations are considered (translation, rotation and reflection). However, in different imaging scenarios it might also become necessary to scale or morph images in order to perfectly align the respective content, *i.e.*, to perform an elastic registration [69, 115, 213]. Generally, a rigid registration of images can be performed either based on intensity value matching or on the registration of extracted features. The former method tries to identify a transformation that maximizes a predefined similarity measure of the intensity values in both images [183].

As this method performs an iterative optimization of the transformation and needs to evaluate the similarity measure at each given step, it is inherently slow and involves the risk of being stuck in a local optimum. Additionally, intensity-based registration heavily relies on matching image content, which might not be satisfied for large developing specimen [184]. Nevertheless, the method can be sped up by restricting the degrees of freedom for the transformation and by choosing a good initialization that might be available from information about the acquisition apparatus. In addition, no landmarks are needed for registration and it is thus possible to perform a markerless multimodal image registration, which is a common task in medical image analysis [140]. Valid similarity measures are, for instance, the sum of squared differences, normalized cross correlation, gradient correlation, difference image entropy or mutual information [175]. A second registration approach is to match extracted landmarks onto each other. These landmarks can for example be scale-invariant features extracted from the actual content of the images [33, 137, 202] or manually placed markers such as fluorescent beads that have to be identifiable in each of the acquired images [106, 185]. Given the landmarks of two complementary views, the registration task comes down to finding an optimal transformation to align the obtained point clouds of landmarks. This approach is computationally much more efficient than intensity-based registration techniques, as the amount of data that has to be aligned is significantly reduced. After having identified corresponding beads in both views, *e.g.*, by local nearest neighbor-based bead descriptors [185] or by approaches like SIFT descriptors [130, 137], the transformation can be calculated iteratively using a random sample consensus (RANSAC) [185] or by direct solutions like least-squares estimations [232].

After having identified an appropriate transformation of the different images, some applications require a fusion of the registered image material in order to obtain a single high-quality image. The simplest method to accomplish a fusion of the images is to perform a local averaging of intensity values. In many cases, however, multiple images are acquired in order to compensate acquisition deficiencies of other views, *i.e.*, it is desirable to only use the high-quality content of the images for the fused result image. More advanced techniques to fuse different images are, *e.g.*, linear blending of the intensity values [230], and entropy-based blending [185] or wavelet-based fusion approaches where a high-quality image is formed by combining the best wavelet coefficients obtained from multiple views [126, 191, 230].

Seed Point Detection

Seed point detection represents an important component in many automated image analysis pipelines. The main intention is to obtain approximate locations of objects of interest, such as pedestrians or vehicles in surveillance videos [14, 82], specific organs or tissues in medical imaging applications [54], particles [71] or labeled nuclei and cells as discussed in this work [1, 155]. Object detection methods are often used as an additional preprocessing step to obtain a precise localization of objects of interest or regions of interest (ROI) that can be further investigated in proceeding analysis steps like seeded watershed, level-sets or graph-cuts segmentations [1, 125, 155, 179]. The major benefit of this preprocessing step is that complex and usually time consuming methods can be intelligently initialized or even guided to the correct image location [1]. In addition to using seed point detection as a preprocessing step, several questions can be reasonably answered using the detected seed points, without the necessity of a subsequent segmentation of the objects of interest. For instance, counting the number of objects, tracking the spatio-temporal dynamics of objects of interest, extraction of registration landmarks or extracting average intensity levels within fixed window are common tasks that can be performed solely based on object centroids and their surrounding [121, 185, 197].

A common method for seed point detection is based on the second-order derivatives of the images using filters such as the Laplacian-of-Gaussian (LoG) filter (Fig. 1.3D) [142] and its approximations Difference-of-Gaussian (DoG)

[185] and Difference-of-Mean (DoM) [77]. These filters efficiently emphasize intensity changes such as edges or bright objects in images of arbitrary dimension at an adjustable scale. The actual seed points are then detected by searching for local extrema in the 8-neighborhood and the 26-neighborhood of the filtered 2D and 3D images, respectively. In order to detect differently sized objects, scale-space approaches use a digitally generated set of successively filtered images to simulate a scaling of the images and thus a scaling of the contained objects [128, 137]. The local extrema detection is then performed within neighboring scales, to determine the scale with the maximum response [137]. Similarly, the wavelet transform can be used to preprocess the images and to detect seeds by multiscale products [168]. Another approach that is frequently used is the concept of Euclidean distance maps (EDM) [145]. In an EDM image (Fig. 1.3F) of a binarized input image (Fig. 1.3E) each foreground pixel value encodes the Euclidean distance to the closest background pixel. Objects of a certain size can be isolated by extracting the h -maxima of the EDM image [210]. A further example for efficient seed detection is the hotspot operator, which performs an analysis of concentric shells surrounding each pixel in an image [242]. Finally, some more complex seed detection approaches are based on partial differential equations (PDE), such as shrinking level-sets [154, 163] or gradient vector diffusion to identify centroids of objects of interest [124]. An exemplary result of a seed detection algorithm is depicted in Fig. 1.3G.

Segmentation

In various application scenarios similar to those listed in the previous section, the localization of an object of interest alone is not sufficient and region boundaries obtained by a segmentation algorithm potentially provide much more information about an object's properties. A detailed analysis of the entire region of an object can be used to extract precise statistical values, regional mean intensities, principal components, covered area and volume, geometrical properties and possible interaction sites between neighboring objects to name but a few. Thus, the goal of the segmentation step is to compute a binary image with a clear and reasonable separation of the objects of interest in the foreground regions of the image (Fig. 1.3E). If multiple objects of interest reside within a single image, it is usually desirable to identify the connected components within the binary image and to assign a unique integer label to each of the individual

objects (Fig. 1.3H). Popular examples for the application of segmentation algorithms in microscopy images are ranging from segmenting fluorescent beads, nano particles or single molecules over labeled cell nuclei and cell membranes through to entire organs or organisms [62, 96, 102, 107, 164, 185].

The easiest way to retrieve a segmented image of bright objects on a dark background is the use of a binary threshold, where all intensity values below the threshold are labeled as background and the remaining pixels get the foreground label assigned. In the case of dark objects on a bright background, the image can be inverted to use the same notion of background and foreground separation. Besides manual adjustment of the thresholds, various methods for the automatic threshold identification have been presented [108, 203]. Most methods for the automatic threshold adjustment are based on an analysis of the histogram mode distribution [169] and can be further improved by additionally considering an image inherent noise model [172]. Khan *et al.* presented a feedback approach to optimize the threshold selection based on annotated ground truth images [100]. However, the application of these straightforward thresholding methods is limited to images with sufficiently high contrast, high signal-to-noise ratio and low object densities and may require pre- or postprocessing steps to improve the result quality. A further frequently used segmentation method is the watershed transform [21, 233], which considers the image intensities as a topographic surface and performs a flooding of the surface starting at local intensity minima. The unseeded version of the algorithm usually strongly over-segments the image and requires additional tricks to correct these flaws [5, 166, 216]. Based on provided seed points of a preprocessing step, segmentation algorithms can be targeted to regions of interest and perform a local extraction of the information present in the images. Examples for seeded segmentation techniques are region growing approaches where seed points are extended based on regional intensity statistics [76], seeded watershed segmentations that start from seeds instead of local intensity minima [15, 222], graph partitioning methods that may use seed points as shape priors [1, 48, 134, 205] or PDE-based methods that can use the provided seed points or seed regions for the initialization of level-set functions or parametric curves [28, 123, 154, 163, 251]. Furthermore, gradient vector flow tracking methods that trace gradient vectors to local minima and use the respective attraction basins as final segmentation [125, 133] or model-based segmentation approaches that try to identify charac-

teristic shapes in the images, for instance, based on elastic registration [52, 95] proofed to provide reasonable segmentation results for particular applications. Recently, classification-based approaches to segmentation have become popular. Based on manually annotated input data, features like image derivatives, intensity distributions or local texture histograms can be used to train random forest classifiers [32, 79, 94, 135] or the features can be even learned automatically from the provided image material using deep learning approaches [49, 89, 109]. The retrieved probability maps of the image classification can then be used to obtain the actual segmentation via thresholding and similar techniques. Obtained segments and the underlying grayscale image regions can be further analyzed, categorized and compared, *e.g.*, using classification approaches as described in [190]. In this work, the focus is put on the segmentation of grayscale images instead of directly analyzing multicolor images in the color space. Nevertheless, it should be noted that many additional algorithms exist for the special task of color image segmentation [207, 240, 254].

Tracking

The tracking of objects is a frequently observed task in spatio-temporal analysis of image data. Generally, the goal of tracking is to determine reasonable one-to-one correspondences in subsequent video frames. Tracking approaches can be found in many different areas, such as surveillance video analysis [29, 212], target tracking [25, 146], pedestrian tracking [178], tracking of virus particles in a host cell [73, 74] or animal tracking for behavioral analyses [31, 223]. In the focus of this work, tracking of thousands of moving cell nuclei was the main application and it was required to quantitatively investigate cellular dynamics in early stages of embryonic development [5, 147, 149, 230]. The following presentation of tracking algorithms thus primarily focuses on valid approaches for the tracking of nuclei and cells in microscopy images.

Depending on the acquisition method, tracking is mostly performed on the results of a preceding segmentation procedure and can either be based on identified segment centroids or on the actual segments. The simplest way to obtain the linkage of the objects in adjoint frames is the identification of the nearest neighbor for each of the contained objects [148]. The performance of these algorithms is reasonably well if the temporal sampling and the segmentation quality is high enough to guarantee unambiguous associations. In addition, these

methods can be easily equipped with various correction heuristics and feature matching approaches to improve the quality [11, 148]. The approach can be further improved using global data association schemes as demonstrated in [23, 93] or global consistency constraints to compensate under-segmentation errors [199]. In addition, some recently presented new tracking approaches identify globally optimal solutions in over-segmented data out of multiple hypotheses [198] or perform the object associations based on structured learning [135].

Another set of tracking algorithms is based on non-parametric and parametric contour evolution [4]. Using an initial shape approximation of the objects, the segmentation masks obtained from approaches like level-sets [28, 127] or parametric contours like Gaussian mixture models [5, 230] are successively propagated to initialize subsequent time points. A related set of methods are state-space models, such as Kalman filters [24, 91] and particle filters [38]. These predictor-corrector schemes estimate future object locations based on a given movement model and the previous object state [4]. A further extension of this approach is given by interacting multiple model (IMM) approaches that additionally allow probabilistic transitions between different movement models based on the given observations [71, 127]. Successfully identified object associations can finally be used to extract spatio-temporal object dynamics, object fate maps and lineage trees in order to visualize and analyze the temporal behavior, ancestry and neighborhood relations of dynamic objects [5, 147, 148, 230].

1.1.3 Benchmarking

To enable a thorough validation of newly developed algorithms and to quantitatively compare different algorithms for a certain task, the extensive use of benchmarking is inevitable. Various benchmarks have already been presented in the past and offer a great toolbox of testing environments for different application scenarios. The benchmarks presented in this section are restricted to the analysis operators relevant to this thesis.

The most straightforward but also most tedious way to obtain a benchmark dataset is the manual annotation of representative image regions. Examples for such benchmarks are manually annotated images for seed point detection [113], for segmentation [51, 70] and for tracking purposes [93, 144]. In order to

circumvent the tedious manual labeling step, many of the recently presented benchmarks rely on simulated data with a directly available ground truth. For instance, in the field of fluorescence microscopy images, various approaches for the simulation of cells and cellular nuclei in 2D and 3D images have been presented [72, 119, 225, 226]. In addition, locally plane-like structures such as cellular membranes can be simulated, *e.g.*, using a Voronoi tessellation of randomly initialized seed points [164]. To additionally add authentic movement behavior of interacting objects, physical phenomena such as repulsion and adhesion interaction can be embedded into the simulations [139]. A great selection of both manually annotated and simulated tracking benchmarks was introduced in the scope of the challenges for particle and cell tracking [45, 144]. Similarly, a set of simulated particle tracking benchmarks can be generated within ICY [41]. A simulated benchmark to assess the quality of multiview deconvolution approaches using spherical objects, point spread function simulations and different noise models is given in [182]. In this work, existing simulated benchmarks were extended, such that entire processing pipelines could be validated on a single comprehensive benchmark.

1.1.4 Prior Knowledge and Uncertainty

The use of prior knowledge is an important component of cutting-edge algorithms. For instance, the approaches described in [1, 28] incorporate information about the number of expected objects and their associated physical size in order to adjust and improve a seed point detection algorithm. Similarly, to perform an image segmentation into foreground and background regions, properties like size, shape, geometry, intensity distributions and the like can be used to improve the algorithmic performance [63, 104, 136]. A frequently used technique to incorporate this prior knowledge is based on shape penalization terms, *e.g.*, by adding additional terms to the energy functional of a graph-cuts [136, 238] or a level-sets segmentation [122] or by generalized Hough transforms that can detect arbitrary shapes [8]. Information about size, shape and movement properties of objects can additionally be used to specify efficient tracking correction heuristics [5, 11, 199].

Besides the use of prior knowledge in image analysis, a great potential lies in the estimation of uncertainties of the automatically produced results. Most

automatic image analysis operators do not produce flawless results and contain a measurable degree of uncertainty that should ideally be considered by subsequent processing steps [196]. A general introduction to the expression of uncertainty inherent to visual sensor signals can be found in [22, 248]. On the pixel level, this uncertainty can be used to assess the information quality of a single pixel due to sensor imperfections or temperature dependence [141, 196]. Furthermore, a lot of previous work was performed to assess the localization uncertainty of geometric features such as corners, centroids, edges and lines in images [6, 42, 43, 173] or to evaluate the quality of image registration algorithms [114]. Besides many applications from the field of quality quantification, a part of research focuses on uncertainty quantification in areas such as face recognition and other biometric technologies [19, 20, 35], the tracking of shapes in ultrasound images [253] or to evaluate the impact of noisy measurements on the validity of diagnosis results [150]. An uncertainty formulation based on fuzzy set theory has been employed to perform pixel-based classification tasks [26, 229] or to detect specific structures in the image [186, 187]. A further possibility to exploit the uncertainty information is to optimize parameter values of a respective operator in a feedback fashion such that the outcome minimizes a previously defined optimization criterion as demonstrated in [100, 101]. A further example is the improvement of a graph-based watershed implementation, where uncertainties are used to assess the influence of individual edges on the final segmentation outcome [222]. In the present thesis, uncertainty is considered as the imperfect knowledge about the validity of a piece of extracted information produced by the respective image analysis operators and was used for efficient improvement heuristics to enhance image analysis pipelines [27, 218].

1.1.5 Available Software Solutions

Of course, a tremendous amount of both commercial and open-source software tools has already been presented to facilitate automatic evaluations of microscopy images and a subset of the most useful ones for analysis tasks similar to those considered in the present thesis is presented here [59]. Many existing tools such as CellProfiler and ZFIQ [37, 133] were specifically developed for 2D applications. Even though they are not directly applicable to 3D analysis

problems, they feature various processing operators that can be arranged in pipelines and executed in batch processing mode, to evaluate large amounts of 2D images and to get an impression of potentially useful algorithms. For both 2D and 3D image analysis tasks, several great open-source software tools have been presented that condense frequently used methods to analyze biological experiments in user-friendly graphical user interfaces (GUI) and are developed and advanced by active communities. Examples for these tools are ImageJ/-Fiji [200], ICY [41], BioImage XD [92] and Vaa3D [174]. A somewhat different approach is pursued by the ilastik application, which essentially uses learning approaches to accomplish image segmentation tasks [211].

Moreover, several high-quality commercial tools like Imaris (Bitplane AG), Velocity 3D (PerkinElmer Inc.), Amira 3D (FEI) and arivis vision (arivis AG) have been developed to facilitate in-depth automated analyses in the life sciences with sophisticated GUIs, special-purpose visualizations and statistical evaluation toolboxes. For rapid prototyping, MATLAB (The MathWorks Inc.) and the associated image processing toolbox represents a valuable tool with an easy-to-learn scripting language, comprehensive image analysis filters and powerful mathematical features for quantitative analyses and visualizations.

To facilitate new implementations, various software development kits (SDK) such as the Insight Toolkit (ITK) [87], the Visualization Toolkit (VTK) [201], OpenCV [30] and the Point Cloud Library (PCL) [194] have been developed and are freely available to the community. Furthermore, BioView 3D [113], VisBio [193] and ParaView [40] are valuable tools for visualization purposes. Some tools that mainly focus on data analysis are, *e.g.*, RapidMiner [153], KNIME [18], WEKA [65] or the MATLAB toolbox Gait-CAD [157]. A comparison of the newly developed XPIWIT software tool to existing software solutions is provided in Tab. 5.1.

1.2 Open Questions

Although previous research in image analysis offers a tremendous amount of elaborate methodology, there are several open questions in the domain of multidimensional image analysis that still are largely unsolved:

- A feasible approach to improve the result quality of automated image analysis operators is the use of prior knowledge. However, available prior knowledge is mostly not sufficiently incorporated to image analysis operators, *i.e.*, a great amount potentially usable extra information is wasted. Furthermore, there is no uniform approach to transform, embed and use the available prior knowledge to improve both existing and new algorithms.
- Although the sequential arrangement of processing operators to image analysis pipelines is a broadly used concept, results propagated through the pipelines are mostly not assessed by the individual pipeline components with respect to their uncertainty. Errors of early processing steps tend to accumulate and may negatively affect the final result quality.
- Existing image analysis methods such as seed point detection and segmentation were mostly developed and tested in 2D scenarios or only work on small image datasets. Thus, they only offer limited applicability to large-scale image analysis tasks and either require to use fiercely resized datasets or immense processing capabilities to obtain quantitative results reliably and fast.
- Many benchmarks for the quality assessment of individual image analysis pipeline components have already been presented. However, comprehensive benchmarks that are crucial to consistently evaluate the quality of all involved steps on the same 3D image data are largely missing. In addition, manually annotated ground truth images may be biased and differently rated by different investigators, *i.e.*, realistic simulated benchmarks are preferable to use in the first place.
- Existing approaches to multiview image fusion are usually based on the assumption that the temporal sampling is sufficiently high or that investigated objects are fixed. If these preconditions are missed, the quality of the fusion images is compromised and may heavily affect the quality of subsequent analysis steps.
- Most available tracking approaches heavily rely on a high-quality segmentation data as their input. However, most methods are not capable of dealing with merged objects resulting in a high number of avoidable linkage errors.

- Even though many usable software tools for multidimensional image analysis have been presented in the past, many applications only offer a limited applicability to terabyte-scale image datasets. Frequently observed problems arise from specific hardware requirements, closed implementations, insufficient parallelization, infeasible memory demands or the inability of being executed in distributed computing environments.
- State-of-the-art microscopy techniques produce a tremendous amount of data with unprecedented quality. However, the investigations and feasible questions have been often compromised by the capabilities of automated processing routines or unattainable infrastructural demands. Hence, many scientific areas demand for sophisticated new approaches to investigate, proof and unveil hypotheses in multidimensional image data reliably, quantitatively and fast.

1.3 Objectives and Thesis Outline

Based on the open problems mentioned in the previous section, the central objectives of the present thesis are:

1. To derive a new concept for assessing and exploiting the uncertainty produced by individual image analysis operators using prior knowledge. The framework should be able to propagate identified uncertainty values through an entire image analysis pipeline and thereby improve the result quality awareness of each processing operator or trigger adapted processing strategies.
2. The development of accurate and fast seed detection and segmentation algorithms that are applicable to detect different geometric shapes in large multidimensional images.
3. New validation benchmarks that closely resemble the real imaging scenarios and allow validating entire processing pipelines with a single consistent dataset.

4. The analysis of seed detection, segmentation, multiview fusion and tracking techniques in different acquisition scenarios and an extension of the algorithms based on the developed uncertainty framework.
5. The implementation of all developed methods in fast, efficient and user-friendly open-source software tools that are applicable to large-scale image analysis problems.
6. The proof of successful operation and performance of all developed methods on large-scale datasets from the field of developmental biology.

To generally improve the incorporation of prior knowledge into image analysis operators and to enable a quantitative assessment of the produced result quality, an extension to the classical image analysis pipeline concept with respect to uncertainty estimation, propagation and usage is presented in Chap. 2. To advance the possibilities of automatic segmentation in terabyte-scale 3D microscopy images, Chap. 3 presents three new segmentation algorithms that were specifically developed to extract predefined geometrical classes of objects and feature highly efficient subroutines, a high degree of parallelization and a low memory footprint to obtain high-quality segmentation results reliably and fast. Chap. 4 illustrates how the uncertainty-based pipeline concept could be used for the enhancement of image analysis components such as seed point detection, segmentation, multiview fusion and tracking and introduces a comprehensive simulated validation benchmark dataset that was suitable for the quantitative assessment of all involved steps. All methods presented in this thesis were implemented as platform-independent open-source software tools as introduced in Chap. 5. In addition to the systematic validation of the new algorithms on synthetic benchmarks, the presented approaches were applied to large-scale microscopy data originating from experiments in developmental biology, which emphasizes the practical applicability of the presented methods (Chap. 6). Finally, a summary of the accomplished tasks and an outlook on potential follow-up projects is provided in Chap. 7.

2 A New Concept for Uncertainty Estimation and Propagation in Image Analysis Pipelines

Many automatically analyzable scientific questions are well-posed and offer a variety of information about the expected outcome *a priori*. Although often being neglected, this prior knowledge can be systematically exploited to make automated analysis operations sensitive to a desired phenomenon or to evaluate extracted content with respect to this prior knowledge. Hence, the performance of processing operators could be greatly enhanced, for instance, by a more focused detection strategy and the direct information about the ambiguity inherent in the extracted data. Using natural language expressions to describe and rate observed phenomena, a new approach to extend image analysis pipelines by the quantification and propagation of uncertainty is introduced in this chapter. Initially, the general image analysis pipeline concept is introduced and an exemplary overview of potentially usable prior knowledge in the domain of microscopy images is provided. To transfer the available background information into a mathematical representation, the concept of fuzzy set membership is presented and it is demonstrated how the obtained feature classification can be used for an improved performance and consciousness of utilized processing operations, *e.g.*, by filtering, propagating and using available information appropriately. The major emphasis is put on the automated analysis of images by means of dedicated processing pipelines. However, the developed concepts can theoretically be applied to similar data processing pipelines observed in other scientific or technical fields with few required adaptations.

2.1 The Image Analysis Pipeline Concept

A general concept inherent to most image analysis pipelines is the use of multiple processing operators that perform a specialized task, *e.g.*, to improve the image signal, to transform the image to a representation that can be more easily processed by downstream processing operators or to extract information from the images (Sec. 1.1.2). The basic layout of such an image analysis pipeline is depicted in Fig. 2.1. Based on a raw input image, the N_{op} sequentially connected operators receive either a processed version of the input image, extracted features or both from their preceding analysis operator. In Fig. 2.1, extracted features that are passed between the pipeline components are denoted by a set \mathcal{X}_i and processed images by \mathbf{I}_i , for the output of the i -th processing operator. More formally, let $i \in \{1, \dots, N_{\text{op}}\}$ be the number of an operator generating a set $\mathcal{X}_i = \{\mathbf{x}_i^\top[n] : n = 1, \dots, N_i\}$ with N_i data tuple vectors containing $N_{f,i}$ features each. The generation of the output set \mathcal{X}_i produced by operator i is directly dependent on the input set \mathcal{X}_{i-1} provided by its predecessor.




Figure 2.1: A common image analysis pipeline comprised of N_{op} processing operators that are sequentially arranged. The performance of each operator is directly dependent on the quality of the input provided by its predecessor. Propagated information can be either images (\mathbf{I}_*), extracted features (\mathcal{X}_*) or both of them (adapted from [219]).

The output set \mathcal{X}_i can also be written as an $(N_i \times N_{f,i})$ matrix:

$$\mathbf{X}_i = \begin{pmatrix} \mathbf{x}_i^\top[1] \\ \vdots \\ \mathbf{x}_i^\top[N_i] \end{pmatrix} = \begin{pmatrix} x_i[1,1] & \cdots & x_i[1,N_{f,i}] \\ \vdots & & \vdots \\ x_i[N_i,1] & \cdots & x_i[N_i,N_{f,i}] \end{pmatrix}. \quad (2.1)$$

The convention used throughout in this chapter is to store each data tuple in a separate row of the table and the individual features to the respective columns. For processing operators such as image preprocessing filters that do not produce any features, \mathbf{X}_i is an empty matrix and only the processed image is passed to the next operator. Of course, image analysis pipelines can be much

more complex than the linear arrangement introduced here. For instance, multiple input images may be given or an analysis operator may require extracted features from multiple sources to generate its results. For convenience, the uncertainty framework introduced in this chapter is illustrated on this basic pipeline structure. The developed concepts, however, are generally applicable and transferring them to more complex pipeline structures is straightforward. An extension of the general result matrix shown in Eq. (2.1) by the quantified object uncertainties derived in Sec. 2.3.1 is provided in Sec. 2.4.

2.2 Identification of Suitable Prior Knowledge

Generally, prior knowledge can be obtained from literature, expert knowledge, experimental evidence or knowledge databases. In the field of microscopy, the instantly available visual representation of the observed data helps to facilitate the identification of recurring patterns, intensity properties or the appearance of objects contained in the images. An exemplary overview of potentially usable prior knowledge from microscopy images is given in the following listing, where italic terms in the brackets correspond to exemplary natural language expressions that could be used to characterize the respective features (adapted from [219]):

- **Microscope:** Acquisition specific prior knowledge such as illumination conditions, detection path, image resolution, physical spacing of voxels, high-quality image regions, point spread function (PSF) or the detection path (*Image quality decreases from XY to YZ*).
- **Intensity:** Signal dependent information like the intensity range, fluorescence properties, signal-to-noise ratio and global statistical properties of the image intensity values (*Valid objects are brighter than XY*).
- **Localization:** Positional information of the objects or object properties in absolute image coordinates. Furthermore, localization of extracted properties or objects relative to each other can be used to define neighborhood relations (*Object type XY only appears close to location YZ*).

- **Spatial Extent:** Object properties such as size, volume, principal components, convex hull extents or bounding volumes (*Object type X is larger than Y but smaller than Z*).
- **Geometry:** Geometrical properties like dimensionality, symmetry, shape, proportions and relative localization of features within an object (*Object type XY has a line-like shape with a central symmetry axis*).
- **Morphology:** Combination of intensity-based and geometrical properties, *e.g.*, to link information about patterning, texture, structure and color to geometrical properties such as shape and symmetry (*Object XY is spherical, bright and has a textured surface*).
- **Object Interaction:** Characterization of between-object properties like clustering, adhesion, repulsion, division or regional density changes (*Object behavior XY rather appears in dense regions*).
- **Spatio-Temporal Coherence:** Dynamically changing quantities such as object growth, movement direction, speed, object appearance and disappearance (*Object moves maximally XY pixels between two subsequent frames*).

The properties are listed in bottom-up order, going from the acquisition stage over the actual content of a single image through to the spatio-temporal comparison between time series of images. Of course, the listing can only provide a coarse overview of potentially usable prior knowledge and the suitable features have to be carefully selected to match the underlying image material and analysis problem. As introduced in the next sections, these linguistic expressions can be used to transform the prior knowledge of different sources to a consistent mathematical representation using the concept of fuzzy sets.

2.3 Prior Knowledge-based Uncertainty Quantification

To fully make use of the prior knowledge presented in the previous section, it has to be transformed into a mathematical representation. A valuable tool for the formulation of prior knowledge in terms of mathematical equations is the use of fuzzy sets that have been introduced by Zadeh in 1965 [252]. Fuzzy sets

extend the concept of classical sets by the incorporation of membership degree, *i.e.*, allowing gradual association of elements to various sets. Linguistic terms are used to describe the fuzzy sets, which makes it a very intuitive approach that is based on human language [10, 105, 110]. Considering the example of gene expression, the use of linguistic expressions enables to quantify vague biological properties like *gene X is slightly over-expressed* and additionally allows soft transitions between expression states instead of using artificially generated sharp thresholds [249].

The concept of fuzzy sets has already been employed in several scientific applications. The technique has been used in the modeling of biological networks and biochemical pathways [2, 116, 161]. Another example is the use of fuzzy clustering with prior knowledge-based distance measures [83, 228]. Moreover, a fuzzy formulation of prior knowledge has been used by Khan *et al.* to automatically adjust parameters of contrast enhancement and segmentation algorithms [100, 101]. A similar formulation of background knowledge using fuzzy set theory is used in this work to assess the uncertainty of information that has been extracted by image analysis operators. An alternative to assess the result quality is, *e.g.*, a Bayesian approach that tries to estimate posterior probabilities for all observed quantities [181]. However, in most cases it is not straightforward to calculate the respective probabilities based on the available data. Soft transitions between feature classifications as well as the possibility to have partial membership associations make fuzzy sets perfectly suited to model the imprecise knowledge observed in many scientific applications [27]. Furthermore, the simplicity of knowledge transformation and the intuitive concepts help, *e.g.*, non-technical users to adjust algorithmic parameters on the foundation of acquired expert knowledge. Based on these benefits the main focus of this work is put on the use of fuzzy sets for the formulation of imprecise information, *i.e.*, imprecisions of available prior knowledge are reflected by the width of the plateau regions with full membership of the respective fuzzy sets. Fuzzy set membership degrees (FSMD) are used to measure uncertainties, which means that uncertainties are considered proportional to the degree of believe a certain statement holds true or a linguistic term is fulfilled [27]. Consequently, high uncertainty corresponds to a low FSMD value to the considered fuzzy set and vice versa, when referring to low uncertainty high FSMD values are assumed.

The fuzzy set framework used in this thesis is mainly employed for a coarse classification of information extracted from images based on vague prior knowledge, in order to decide if produced results match the expectation or may require further processing to correct flaws. In addition, the assessed uncertainty in form of FSMD values of extracted information to predefined fuzzy sets is shared between multiple pipeline components to allow more elaborate decisions in each step. It should be noted, though, that the use of the term uncertainty propagation does not refer to the statistical error propagation theory but only indicates that estimated FSMD values are shared between operators of a processing pipeline.

2.3.1 Quantifying Prior Knowledge using Fuzzy Set Membership Functions

In the classical sense, a set \mathcal{C} is characterized by clearly defined boundaries that can be described with its associated characteristic function [27]:

$$\chi_{\mathcal{C}}(x) = \begin{cases} 1, & x \in \mathcal{C} \\ 0, & x \notin \mathcal{C}, \end{cases} \quad (2.2)$$

which means that an element can either be part of the set or not. An extension to this definition of a set are so called fuzzy sets, which allow a gradual membership of elements. Analogous to the characteristic function of a classical set, a fuzzy set \mathcal{A} can be defined by its associated membership function (MBF) $\mu_{\mathcal{A}} : \mathcal{X} \mapsto [0, 1]$ that maps each element of a universe of discourse \mathcal{X} to a value in the range $[0, 1]$ [252]. This assigned value in turn directly reflects the degree of membership (FSMD) of the respective element to the fuzzy set \mathcal{A} . In the special case $\mu_{\mathcal{A}}(x) = 1$, the element x is fully included in \mathcal{A} , whereas $\mu_{\mathcal{A}}(x) = 0$ indicates that x is not part of the fuzzy set \mathcal{A} at all [16]. Examples for possible membership functions are singletons, triangular, rectangular, trapezoidal, sigmoidal and Gaussian MBFs (Fig. 2.2). The most common membership functions used in practice are trapezoidal membership functions, which can additionally be parameterized to model singletons, triangular and rectangular MBFs [105].

A trapezoidal membership function can be formulated as:

$$\mu_{\mathcal{A}}(x, \boldsymbol{\theta}) = \max \left(\min \left(\frac{x-a}{b-a}, 1, \frac{d-x}{d-c} \right), 0 \right), \quad (2.3)$$

with the parameter vector $\boldsymbol{\theta} = (a, b, c, d)^\top$ that is used to control the start and end points of the respective transition regions. An appropriate MBF type has to be selected on the basis of the underlying problem and mostly the use of a standard partition represents a reasonable choice [57, 110]. In a standard partition, maximally two neighboring fuzzy sets overlap and have non-zero membership values for a certain value of x . At the same time, the membership degrees for any input value of x to the corresponding fuzzy sets have to sum up exactly to 1 with plateau regions indicating that the respective element is fully included in the fuzzy set \mathcal{A} (Fig. 2.3).




Figure 2.2: Exemplary membership functions that can be used to define fuzzy set membership degrees. The panels show a trapezoidal (A), a triangular (B), a singleton (C), a sigmoidal (D) and a Gaussian membership function (E). The triangular and singleton membership functions are special cases of the trapezoidal membership function and all of them can be parameterized using the parameters a, b, c, d that define the locations of the respective transition regions (A).

The soft transitions of an object's membership between different fuzzy sets are perfectly applicable to model imprecise experimental observations, *e.g.*, obtained from biological experiments [27]. In particular, a fuzzy set \mathcal{A} can be used to express the gradual validity of a linguistic term describing a certain phenomenon by one or more feature values. Considering, for instance, the detection of a specific kind of object it might make sense to use the size feature of the object as an indicator of its appropriateness. Linguistic terms could in this case for example be restricted to five possible outcomes, where the extracted feature ...

1. ... perfectly matches the expected value (*Correct*).
2. ... is smaller than expected but might be useful (*Small*).
3. ... is larger than expected but might contain useful information (*Large*).
4. ... is too small and not useful (*Too Small*, e.g., noise or artifacts).
5. ... is too large and not useful (*Too Large*, e.g., segments in background regions).

The parameterization of the associated fuzzy sets can be performed on the basis of available prior knowledge and an exemplary standard partition for these five linguistic terms using trapezoidal membership functions is shown in Fig. 2.3A.




Figure 2.3: Different possibilities to partition the input space of a feature x using trapezoidal membership functions. In (A), each of the linguistic terms has a separate fuzzy set and (B) shows a reduced version with only three fuzzy sets that correspond to the desired class and its complement. In (B), different possibilities to summarize the correct objects arise. Besides solely restricting the class to the correct set as done in (B), the correct fuzzy set could be extended by the potentially useful classes (*Small* and *Large*). However, the appropriate formulation has to be chosen application dependent (adapted from [219]).

Considering an entire processing pipeline that may be comprised of N_{op} processing operators, it is theoretically possible to specify multiple linguistic terms for each extracted feature of each processor $i \in \{1, \dots, N_{\text{op}}\}$ if the respective prior knowledge is available and reasonably usable. Put more generally, a feature $f \in \{1, \dots, N_{f,i}\}$ produced by the image analysis operator i is fuzzified by N_l linguistic terms, which may be specified based on *a priori* knowledge. The fuzzy set membership functions for image analysis operator i , feature f and linguistic term $l \in \{1, \dots, N_l\}$ are then denoted by $\mu_{A_{ifl}} : \mathbb{R} \rightarrow [0, 1]$ (Fig. 2.3A).

A data tuple $\mathbf{x}_i[n]$ obtains the degree of membership $\mu_{\mathcal{A}_{i,fl}}(x_i[n, f])$ to the fuzzy set $\mathcal{A}_{i,fl}$ for each feature f and each linguistic term l .

If only one outcome of the operators is of importance (e.g., case 1 in the above-mentioned example), it is possible to use only one linguistic term and aggregate all other cases by its complement, which is given by $\mu_{\mathcal{A}_{i,f2}}(x) = 1 - \mu_{\mathcal{A}_{i,f1}}(x)$ (Fig. 2.3B). Of course, this formulation only allows to determine if an object belongs to case 1 or not. However, depending on the field of application it might also make sense to explicitly use the remaining cases, too. Considering for instance an image segmentation operator that erroneously extracts fused image regions due to under-segmentation. Although the information might not directly be useful, the information that some of the objects are classified to case 3, i.e., having a large region that potentially contains desired objects, might become important at later processing stages (Sec. 2.4). To facilitate the implementation this can be achieved by using only a single fuzzy set to determine the objects that need to be corrected (low FSMD values to case 1) and subsequently determine if they are too large or too small by a simple comparison to the boundaries of the *Correct* fuzzy set.




Figure 2.4: Color-based visualization of the fuzzy set membership degree for a seed detection algorithm (A) and a segmentation algorithm (B). Red objects have low membership values due to being noise or merged regions, blue objects are potentially useful but may contain special object properties and green objects perfectly match the expectation (adapted from [218]).

Fig. 2.4 shows exemplary results of two image analysis operators that were applied on a microscopy images showing stained nuclei of HeLa cells (Raw data taken from http://www.cbi-tmhs.org/Dcelliq/files/051606_

HeLaMCF10A_DMSO_1.rar). The degree of membership to the desired class of an appropriately detected cell is indicated by the color code following the fuzzy sets depicted in Fig. 2.3B. Objects with size feature values lying in the expected range are colored in green, potentially useful objects that are slightly smaller or larger than expected are colored in blue and finally noise segments as well as merged segments are colored in red. The five different fuzzy sets used to characterize the objects by their size (Fig. 2.3A) can be exploited, *e.g.*, to split merged objects that are considered too large (case 3) or to filter out false positive detections in background regions (case 4). The next section describes how the membership values of different features can be combined to form more complex object descriptors.

2.3.2 Combination of Fuzzy Set Membership Functions

In many cases a specific phenomenon and the associated linguistic term are not only characterized by a single feature but by a combination of concurrently valid assumptions. Therefore, it is desirable to be able to combine the individually specified fuzzy sets to a single expression that yields a combined membership function for a particular linguistic term. For instance, if an extracted object has to be structured and rectangular to be a desired object, the corresponding FSM values obtained for both features should be high, *i.e.*, they should be combined using a fuzzy pendant to a logical conjunction [27]. A conjunction of $N_{f,i}$ fuzzy set membership functions for linguistic term l based on the multiplication operator can be formulated as:

$$\mu_{lc, \mathcal{A}_{il}}(\mathbf{x}_i[n]) = \prod_{f=1}^{N_{f,i}} \mu_{A_{ifl}}(x_i[n, f]), \quad (2.4)$$

where the index of μ_{lc} simply indicates that the membership function is a combination of multiple linguistic terms. Similarly, a conjunction can be defined using the minimum operator:

$$\mu_{lc, \mathcal{A}_{il}}(\mathbf{x}_i[n]) = \min_{f=1, \dots, N_{f,i}} (\mu_{A_{ifl}}(x_i[n, f])). \quad (2.5)$$

Features that should be excluded from the multiplication (Eq. (2.4)) or the minimum calculation (Eq. (2.5)) can be disabled by setting the corresponding MBFs

to the constant value 1 (identity element of the conjunction). The complement of the combined linguistic term is simply given by $1 - \mu_{lc, \mathcal{A}_{ii}}(\mathbf{x}_i[n])$ as illustrated in Fig. 2.3B. In both formulations, high values of the combined functions are obtained only if all individual feature fuzzy sets yielded high values, *i.e.*, a low value of the combined MBFs indicates that the investigated object strongly deviates from the expected properties of the respective linguistic term. The minimum-based formulation in Eq. (2.5) is somewhat more informative than the multiplication, as a non-zero value directly represents the lowest FSMD value of the considered fuzzy sets. Of course, other logical combinations like disjunctions or implications are in principle possible but the conjunction expressions mentioned above were sufficient for the applications of this thesis.

2.4 Uncertainty Propagation in Image Analysis Pipelines

In the previous section it has been shown, how prior knowledge can be quantified using fuzzy set theory to characterize and assess the extracted information of a single processing operator according to defined linguistic terms. However, image analysis pipelines are mostly comprised of multiple processing operators that are serially linked and directly dependent on the results provided by previous calculations. Thus, an important extension is to pass the quantified uncertainty associated with each data tuple through the pipeline, *i.e.*, performing a feed-forward propagation of the reliability of data shared between operators. For each data vector that is produced by operator i , the degree of membership to the respective fuzzy sets is calculated and the initial feature output matrix defined in Eq. (2.1) is replaced by an augmented matrix of size $(N_i \times (N_{f,i} + N_1))$ that is passed to the subsequent operators for further processing steps:

$$\mathbf{X}_i = \begin{pmatrix} x_i[1, 1] & \cdots & x_i[1, N_{f,i}] & \mu_{lc, \mathcal{A}_{i1}}(\mathbf{x}_i[1]) & \cdots & \mu_{lc, \mathcal{A}_{iN_1}}(\mathbf{x}_i[1]) \\ \vdots & & \vdots & \vdots & & \vdots \\ x_i[N_i, 1] & \cdots & x_i[N_i, N_{f,i}] & \mu_{lc, \mathcal{A}_{i1}}(\mathbf{x}_i[N_i]) & \cdots & \mu_{lc, \mathcal{A}_{iN_1}}(\mathbf{x}_i[N_i]) \end{pmatrix}. \quad (2.6)$$

In Eq. (2.6), it is theoretically allowed to add an arbitrary number of combined linguistic terms N_l of additional criteria $\mu_{l_c, \mathcal{A}_{il}}(\mathbf{x}_i[n])$. If only a classification into correct/incorrect objects needs to be performed, one linguistic term is sufficient ($N_l = 1$). Note that the combined linguistic terms μ_{l_c} obtain the entire feature vector $\mathbf{x}_i[n]$ of data tuple n as an argument. This notation is chosen for brevity and it does not necessarily mean that each entry of $\mathbf{x}_i[n]$ contributes to the combined fuzzy set membership function. If the conjunction operator is used for the combined linguistic terms, the fuzzy sets of features that should not contribute to the combined linguistic term are set to a constant value of 1.

2.4.1 Uncertainty-based Object Rejection

The first application of the uncertainty framework is to filter the extracted output information \mathcal{X}_i produced by an operator i using thresholds $\alpha_{il} \in [0, 1]$. Each processing operator i that is part of the pipeline obtains an individual threshold parameter α_{il} for each linguistic term l , e.g., parameter α_{31} would be used to threshold the FSMD value of the first linguistic term of the third processing operator. According to the FSMD values $\mu_{l_c, \mathcal{A}_{il}}(\mathbf{x}_i[n])$ calculated for each data tuple $\mathbf{x}_i[n] \in \mathcal{X}_i$, $\mathbf{x}_i[n]$ is only passed to the next pipeline component if $\mu_{l_c, \mathcal{A}_{il}}(\mathbf{x}_i[n]) \geq \alpha_{il}$ for the membership to a desired set. The reduced set which serves as input for operator $i + 1$ can be written as:

$$\tilde{\mathcal{X}}_i := \{\mathbf{x}_i[n] \in \mathcal{X}_i : \mu_{l_c, \mathcal{A}_{il}}(\mathbf{x}_i[n]) \geq \alpha_{il}, \forall l \in \{1, \dots, N_l\}\}. \quad (2.7)$$

This object suppression based on the respective FSMD values allows to easily filter out false positive detections based on predefined criteria of the expected object properties. For instance, to improve the sensitivity of an object detection algorithm it might be necessary to use parameters that produce many false positive detections in background regions but offer a high detection rate. Prior knowledge can be used to obtain a coarse classification of the objects and to reject inappropriate objects from further processing. Examples for the application of this filtering approach to improve the precision of a seed point detection and a segmentation algorithm are provided in Sec. 4.2 and Sec. 4.3, respectively.

2.4.2 Extended Information Propagation to Compensate Operator Flaws

In a usual image analysis pipeline, each processing operator mostly only uses the input of its direct predecessor. However, in some cases it might be advantageous to the processing operators to be able to access results obtained by earlier processing steps. For instance, if an operator i fails to sufficiently extract information from its provided input data (*e.g.*, missing, merged or misshapen objects), it can inform downstream operators about these flawed results. The characterization of the produced results is again performed on the basis of FSMMD values to the prior knowledge-based fuzzy sets obtained by the respective extracted objects. For extracted objects that exhibit a low degree of membership to the desired set, it might be beneficial to keep information $\mathbf{x}_{i-1}[n] \in \tilde{\mathcal{X}}_{i-1}$ from one of the previous steps that produced the uncertain information $\mathbf{x}_i[n] \in \mathcal{X}_i$. Using a second threshold $\beta_{il} \in [\alpha_{il}, 1]$ for each operator, the FSMMD level below which the information of the previous steps should be additionally propagated can be controlled. Analogously to the forward threshold, each processing operator i that is part of the pipeline obtains an individual threshold parameter β_{il} for each linguistic term l , *e.g.*, parameter β_{42} would be used to propagate information based on the FSMMD value of the second linguistic term of the fourth processing operator. More formally this means that instead of only forwarding the α_{il} -filtered set $\tilde{\mathcal{X}}_i \subseteq \mathcal{X}_i$ produced by operator i to operator $i + 1$ a set

$$\Omega_i = \tilde{\mathcal{X}}_i \cup \tilde{\Omega}_{i-1}, \quad (2.8)$$

with $i \geq 2$ and $\Omega_1 = \tilde{\mathcal{X}}_1$ is passed through the pipeline. In Eq. (2.8), $\tilde{\Omega}_{i-1}$ is defined as:

$$\tilde{\Omega}_{i-1} = \{\mathbf{x}_{i-1}[n] \in \Omega_{i-1} : \alpha_{il} \leq \mu_{lc, \mathcal{A}_{il}}(\mathbf{x}_i[n]) < \beta_{il}, \mathbf{x}_i[n] \in \tilde{\mathcal{X}}_i, \forall l \in \{1, \dots, N_l\}\}, \quad (2.9)$$

and thus represents the subset of elements in Ω_{i-1} which were not successfully transferred into useful information by operator i , *i.e.*, elements $\mathbf{x}_{i-1}[n] \in \Omega_{i-1}$ that generated output $\mathbf{x}_i[n] \in \tilde{\mathcal{X}}_i$ with $\alpha_{il} \leq \mu_{lc, \mathcal{A}_{il}}(\mathbf{x}_i[n]) < \beta_{il}$. Such elements characterize information of operator $i - 1$ that might be useful in later steps to correct flawed results of operator i . Although there is in principle no limitation on how many predecessors provide their information in a feed-forward fash-

ion, it mostly makes sense to only use a part of the available data depending on the information that is actually needed by downstream operators to avoid unnecessary communication overhead. In the current version of the framework it was assumed to be the responsibility of the respective processing operators to calculate Ω_i and to appropriately calculate the respective FSMD values. This approach was successfully applied to improve the outcome of a tracking algorithm that was allowed to use both extracted segmentation masks and seed points that produced the segments. In cases where the provided segmentation data was insufficient, the algorithm could fall back on seed point information of the penultimate processing operator to resolve tracking conflicts that originated from under-segmentation errors as described in Sec. 4.5.

2.4.3 Resolve Ambiguities using Propagated Uncertainty

Next to filtering and propagating the information of the operators within the pipeline, the operators can explicitly make use of the provided uncertainty information to improve their results. Depending on the degree of uncertainty of provided information, parameters or even whole processing methods might be adapted if needed. Of course, the adaptations required by a particular algorithm cannot be generalized and only exemplary applications are introduced in this section. An important potential application is the correction of under-segmentation errors. The evaluated FSMD values can be used to identify segments that deviate from the expectation (larger than expected, corresponding to case 3 in Sec. 2.3.1) and the identified segments can then be separated using additional information such as seed points or distance maps. An exemplary segmentation algorithm based on adaptive thresholding was successfully equipped with the proposed extension as demonstrated in Sec. 4.3.

Another promising application is the fusion of redundant information that might originate from different sensors, multichannel or multiview experiments or to combine the results of different algorithms applied on the same dataset. After having identified a likely grouping of the data using common classification or clustering approaches, the combination of redundant information within a class or cluster can be performed using the propagated FSMD values. Possible ways to combine the redundant information based on the respective FSMD values are a convex combination of the feature vectors with weights

proportional to the FSMD values (*e.g.*, for the fusion of redundant seed points) or to simply select the object that yielded the highest FSMD value (*e.g.*, for multiview segmentation fusion). Using this approach, it was possible to fuse redundant data extracted from complementary multiview images as shown in Sec. 4.4 and Sec. 6.2.

2.4.4 Improved Processing Pipeline by Uncertainty Propagation

The general scheme for the proposed uncertainty propagation framework is summarized in Fig. 2.5. All in all, the uncertainty framework described in this chapter encompasses filtering, propagation and exploitation of information produced by and between individual pipeline components based on fuzzy set membership functions.




Figure 2.5: In most analysis pipelines each of the N_{op} processing objects directly relies on the output \mathcal{X}_i of its predecessor. Using thresholds the extracted output information of each operator can be filtered according to its uncertainty leading to a reduced set $\tilde{\mathcal{X}}_i \subseteq \mathcal{X}_i$ that is passed to the successor ($*_1$, threshold α_{il}). If an operator i produces uncertain outcome $\mathbf{x}_i[m] \in \mathcal{X}_i$ its successor may access preserved information from one of the previous steps. Therefore, each element $\mathbf{x}_{i-1}[m] \in \Omega_{i-1}$ that led to an uncertain result in processing object i is kept in the set $\tilde{\Omega}_{i-1} \subseteq \Omega_{i-1}$ and additionally passed to the successor, *i.e.*, a subset of $\Omega_i = \tilde{\mathcal{X}}_i \cup \tilde{\Omega}_{i-1}$ with $i \geq 2$ and $\Omega_1 = \tilde{\mathcal{X}}_1$ is used as input for operator $i + 2$ ($*_2$, threshold β_{il}). Solid lines indicate the main information flow, dash-dotted lines the propagation of previously calculated results and dotted lines emphasize influence on the selection of propagated information. In addition to the flow of extracted features (\mathcal{X}_* , Ω_* , $\tilde{\mathcal{X}}_*$, $\tilde{\Omega}_*$), the operators may pass processed images, image parts or simply forward the input image (\mathbf{I}_*) to the subsequent processing operator (adapted from [219]).

The key ideas of the proposed framework are summarized in the following enumeration:

1. Filter information based on their uncertainty, *i.e.*, define uncertainty thresholds α_{il} for each pipeline component. To keep all extracted information the threshold is set to $\alpha_{il} = 0$. Contrary, for $\alpha_{il} = 1$, no uncertain information is passed to operator $i + 1$ (Sec. 2.4.1).
2. Allow operators to fall back on information of penultimate processing steps if predecessors do not deliver good results. The parameter β_{il} specifies the level up to which operator i should preserve previous information. If $\beta_{il} = 1$ all information of $i - 1$ that produced an uncertain outcome is propagated to the successor $i + 1$. If $\beta_{il} = \alpha_{il}$ only the information $\tilde{\mathcal{X}}_i$ produced by operator i is propagated (Sec. 2.4.2).
3. Specifically handle information provided by previous operators based on the respective uncertainty, *e.g.*, use the β_{il} -threshold to inform an algorithm which of its produced results might need to be revised or to resolve ambiguities when combining redundant information (Sec. 2.4.3).

The proposed uncertainty framework is exemplarily applied to some of the newly developed segmentation approaches of Chap. 3. As described in Chap. 4, an existing image analysis pipeline comprised of seed point detection, segmentation, multiview fusion and object tracking greatly benefited from the uncertainty-based improvements and the extensions only had a low impact on the processing times due to the lightweight adaptations of the propagated feature matrices.

3 New Approaches for Efficient Segmentation in Multidimensional Image Data

In the previous chapter, the estimation and propagation of prior knowledge-based validity of extracted information was presented as a general technique to enhance arbitrary image analysis operators. Here, three newly developed segmentation algorithms are introduced, which on the one hand serve as exemplary processing operators to flesh out the general concepts presented in the last chapter and on the other hand represent efficient new solutions to three important problems observed in large-scale image analysis.

Arguably, image segmentation is one of the most important processing steps in numerous practical applications that basically tries to determine an image partition into a finite set of regions. In many cases the segmentation task consists of extracting objects that belong to a specific class of geometry and one way to achieve this segmentation goal is the explicit incorporation of prior knowledge into the developed algorithms. Examples for algorithms that can be guided to detect specific predefined shapes are the Hough-transform [8], shape priors for graph-cut and level-set segmentation approaches [66, 162] or eigenvalue-based approaches that emphasize a specific type of geometry [7, 64]. Existing segmentation algorithms, however, were mostly developed for 2D applications and many mathematically elegant concepts and implementations were optimized with a focus on segmentation quality rather than computational efficiency. On the other hand, computationally efficient methods such as adaptive thresholding or blob detection methods mostly disregard prior knowledge and require sophisticated pre- and postprocessing procedures to correct erroneous results. To enable analyses of the tremendous amount of image data produced by state-of-the-art multidimensional imaging techniques, however, particularly the computational efficiency is a mandatory criterion that has to be met while at the same time providing high segmentation quality. Another common issue

is the lack of thorough validation, *i.e.*, the performance of new methods is often only tested on a particular type of data without consideration of possible image flaws like noise or out-of-focus effects.

The approaches presented in this chapter were specifically designed with these limitations of existing algorithms in mind. Each algorithm was tailored to detect a specific geometric class of objects and thereby maximally exploits available prior knowledge to obtain accurate segmentation results. First, the problem was analyzed thoroughly and it was generalized to an abstract geometrical class. To be able to assess the performance of both existing and new algorithms under various image quality conditions, simulated benchmarks were used in combination with different image distortion techniques. This approach of extensive benchmarking throws light on the particular strengths and weaknesses of the algorithms and helps to narrow down reusable existing components and unsuitable components that require improvements with respect to time performance or segmentation quality. Parameterization of existing methods should of course be performed with the maximal amount of available prior knowledge to guarantee their optimal operation. Performance improvements were achieved by incorporating efficient processing operators and by heavily parallelizing the calculations where possible, to make large-scale analyses feasible. Developed pipelines were then further optimized by the explicit consideration of available prior knowledge in postprocessing heuristics to filter and improve extracted information (Chap. 4). Finally, the pipelines were adapted and optimized to the particular application they were developed for (Chap. 6). The general concept that was used for the algorithmic development of the new segmentation methods is sketched in Fig. 3.1.

The first project considered in this thesis required a detection and extraction of the spinal cord in laterally oriented zebrafish embryos, *i.e.*, more generally the extraction of elongated, possibly curved structures in 2D images including a characterization of the local environment (Sec. 3.2, [221]). Besides compensating the usual deficiencies of fluorescence microscopy images, the algorithm needed to be able to accurately extract comparable regions of interest from possibly thousands of 2D images on a single desktop computer. Furthermore, an efficient pipeline for the detection, segmentation and tracking of fluorescently labeled cellular nuclei and membranes in large-scale 3D+t microscopy images of various model organisms had to be developed. To be able to even analyze

multiple terabytes of 3D+t image data, the algorithms had to be both highly efficient and reasonably accurate to validate biological hypotheses. Generally, these developed methods are suitable for a fast and accurate segmentation of spherical objects (Sec. 3.3, [220]) and locally plane-like objects (Sec. 3.4, [216]) in large-scale 3D images. Although the algorithms were developed with these particular applications in mind, the thorough validation on more general synthetic benchmark datasets confirms their usability and robustness, and the general concepts can be easily adapted to similar problems in other scientific fields. As all presented segmentation approaches make use of a seed detection stage for an approximated object localization, two seed detection algorithms that are well suited for this purpose are presented in Sec. 3.1.




Figure 3.1: Schematic overview of a general procedure for the development of a new segmentation algorithm. Initially, the problem is analyzed and abstracted to create a realistic simulated benchmark. Flaws and benefits of existing methods are systematically identified based on the benchmark and prior knowledge-based parameter optimization. A new algorithm or pipeline is then developed using existing components where possible and new components where computational efficiency or quality improvements are required. Available prior knowledge is exploited to generate sophisticated postprocessing heuristics and finally the developed pipeline is adapted to the actual application scenario.

3.1 Seed Point Detection

In Sec. 1.1.2, the conceptual idea of a seed detection stage in the image analysis pipeline was presented. Although the projects considered in this thesis specifically required to extract fluorescently labeled nuclei and cells in 3D+t microscopy images, the described methods can in principal be used to detect

(hyper-)spherical objects in arbitrary dimensions. Here, a blob detection using Laplacian-of-Gaussian (LoG) filtering [12, 220] and a second approach based on the Euclidean distance transform [216] are described. Although both approaches are established methods for seed point detection, some modifications were necessary to improve their applicability to large-scale datasets and to enhance their detection capabilities.

3.1.1 Validation Benchmark

The validation of the presented seed detection algorithms was performed on a set of 3D benchmark images that contained simulated spherical objects as described by Svoboda *et al.* and were generated using the CytoPacq simulation toolbox (HL60 cell line, low SNR, 75% clustering probability, http://datasets.gryf.fi.muni.cz/cytometry2009/HL60_HighNoise_C75_3D_TIFF.zip) [225].




Figure 3.2: Maximum intensity projections of differently flawed benchmark images with different levels of additive Gaussian noise (A-D) and Gaussian smoothed versions with fixed noise level (E-H). While the low noise levels in (B) can usually be handled by seed detection algorithms, higher noise levels (C, D) cause a lower signal-to-noise ratio and thus the structures of interest become eventually almost indistinguishable from the background noise. Increasing the variance of the Gaussian smoothing causes initially separated objects to merge and sharp edges of the objects are successively blurred.

To additionally assess the performance of the seed detection algorithms in relation to the signal-to-noise ratio, a series of images with increasing zero-mean additive Gaussian noise levels with a standard deviation σ_{agn} in the range of $[0.0, 0.128]$ have been used. Furthermore, out-of-focus effects were simulated

using Gaussian blurred versions of the benchmark images with different variances σ_{smooth}^2 in the range of $[0.0, 100.0]$ and a fixed zero-mean additive Gaussian noise with $\sigma_{\text{agn}} = 0.001$. The ranges of both image distortion approaches were empirically determined, where the largest values produced images with signal-to-noise ratios close to the visual detection limit [165]. The specifiers SBDS1, SBDS2 and SBDS3 are used to refer to the raw, the noisy and the smoothed datasets of the benchmark, respectively. Exemplary benchmark images are shown in Fig. 3.2 and the settings of all datasets are summarized in App., Tab. D.1.

3.1.2 Seed Detection using a Laplacian-of-Gaussian Maximum Projection

The Laplacian-of-Gaussian filter and its approximations are well suited for the detection of spherical objects in images of arbitrary dimension (Sec. 1.1.2). To be able to detect differently sized objects, LoG-filtered images at different scales were used [1, 128] in combination with efficient recursive Gaussian filtering approaches that enabled seed detection even on large volumetric images [56]. The LoG-filters were scale-normalized, *i.e.*, the 3D filtering considered the physical spacing of the voxels to compensate the reduced axial resolution which is common in 3D microscopy and tomography approaches:

$$\mathbf{I}^{\text{LoG}}(\mathbf{x}, \sigma) = \mathbf{I}^{\text{raw}}(\mathbf{x}) * \left(\sigma^3 \sum_{i=1}^3 \frac{\partial^2 \mathbf{G}(\mathbf{x}, \sigma)}{\partial x_i^2} \right), \quad (3.1)$$

with \mathbf{I}^{raw} being the input image, $\mathbf{x} = (x_1, x_2, x_3)^\top$ being the voxel coordinates, $\mathbf{G}(\mathbf{x}, \sigma)$ representing a Gaussian kernel with standard deviation σ and $*$ being the convolution operator. For performance reasons, however, the automatic scale selection performed in [1] was skipped and the calculations were restricted to a set of scales that was determined based on the radius of the objects of interest. Using the relationship that the radius of detected objects is $r = \sqrt{2} \cdot \sigma$, the appropriate set of scales could be determined by *a priori* knowledge about the investigated biological content of the images [1]. Using the radius of the smallest object of interest r_{min} yields $\sigma_{\text{min}} = r_{\text{min}}/\sqrt{2}$. Analogously, the radius of the largest object of interest r_{max} can be used to obtain $\sigma_{\text{max}} = r_{\text{max}}/\sqrt{2}$. If σ_{min} is too small ($\sigma_{\text{min}} \ll r_{\text{min}}/\sqrt{2}$), an increased amount of

false positives is detected. Analogously, if the parameter σ_{\max} is set too large ($\sigma_{\max} \gg r_{\max}/\sqrt{2}$), nearby objects may become indiscernible, which potentially causes missed detections. In most common scale-space approaches, points of interest are detected in the local surrounding of the investigated point within neighboring scales and possibly even at different octaves [137]. Due to the enormous image size of several gigabytes confronted in the discussed application example, it was not feasible to keep multiple image stacks of the scale-space simultaneously in memory. Therefore, an iteratively calculated LoG scale-space maximum projection with a predefined discrete step size was used instead:

$$\mathbf{I}^{\text{LoGMP}}(\mathbf{x}, \sigma_{\min}, \sigma_{\max}) = \max_{\sigma_{\min} \leq \sigma \leq \sigma_{\max}} \mathbf{I}^{\text{LoG}}(\mathbf{x}, \sigma). \quad (3.2)$$

For each pixel the scale that yielded the maximum value was additionally stored:

$$\mathbf{I}^{\text{MS}}(\mathbf{x}, \sigma_{\min}, \sigma_{\max}) = \operatorname{argmax}_{\sigma_{\min} \leq \sigma \leq \sigma_{\max}} \mathbf{I}^{\text{LoG}}(\mathbf{x}, \sigma). \quad (3.3)$$

This information could be exploited for down-stream processing operators as it provided an initial size approximation of the object about to be extracted. Finally, the actual seed extraction from the LoG scale-space maximum projection image came down to a simple local extrema detection in the direct neighborhood of each pixel (8-neighborhood and 26-neighborhood for 2D and 3D images, respectively). To discard detections in background regions, an intensity threshold was used to only consider seed points whose intensities were above the specified level. The intensity threshold can either be determined automatically, *e.g.*, using Otsu's method [169] or via manual adjustment as described in Sec. 5.2.6. Instead of only using the single pixel value at the seed location, the robustness of the intensity-based filtering approach could be increased by extracting the average intensity in a small window centered at the seed location and by performing the threshold operation on this average intensity value. An efficient way to estimate the window mean intensity at a seed point location was to extract the intensity values from an integral image-based mean filtered raw image [234]. A schematic overview of the introduced pipeline is depicted in Fig. 3.3.




Figure 3.3: Pipeline overview of the LoG-based seed detection. Local extrema are extracted from a LoG scale-space maximum projection. Additional features such as the mean intensity in the vicinity of each local extremum can be used to further filter extracted seed points, *e.g.*, to suppress false positive detections. Processing steps shaded in grey incorporate prior knowledge such as expected radius and intensity properties of the desired objects. The result filter is indicated by a dashed edge line.

The processing steps needed to obtain the LoG scale-space maximum projection for an exemplary 2D image and an overlay of the detected seed points on the raw input image are shown in Fig. 3.4. In Sec. 4.2.1 and Sec. 4.2.2, respectively, an extension of the described LoG-based seed detection algorithm with respect to object fusion and uncertainty-based object rejection is introduced, which significantly improved the detection quality and additionally provided information about the validity of extracted seed points.




Figure 3.4: Visualization of the processing steps that are required to obtain a Laplacian-of-Gaussian (LoG) scale-space maximum projection for seed detection. (A) Original image showing spherical objects of the SBDS1 benchmark images, (B, C) LoG filtered image ($\sigma = 10$, $\sigma = 13$), (D) LoG scale-space maximum projection ($\sigma_{\min} = 10$, $\sigma_{\max} = 13$, $\sigma_{\text{step}} = 1$) and (E) a superposition of detected seeds on the raw image. As all contained objects had a similar radius, a single scale would have been sufficient ($\sigma = 10$) and the additional scale as well as the maximum projection are only visualized for illustration purposes (adapted from [220]).

3.1.3 Seed Detection using Thresholding and Euclidean Distance Maps

A second approach to detect seed points for (hyper-)spherical objects like fluorescently labeled nuclei in multidimensional images is the use of Euclidean distance maps (EDM) of a binarized input image and is schematically illustrated in Fig. 3.5.

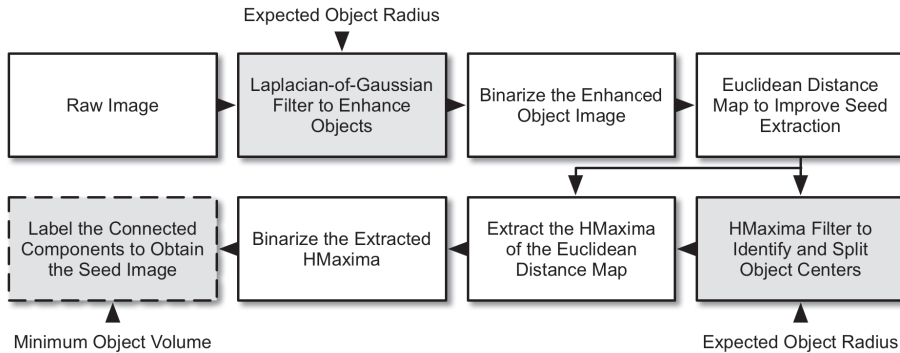


Figure 3.5: Schematic pipeline overview of the EDM-based seed detection. Seeds are extracted from the distance map of a preprocessed, binarized object image. Operators shaded in gray incorporate prior knowledge about minimum and maximum expected object radii and volume to improve the detection of desired objects and to discard noise detections. The result filter is indicated by a dashed edge line.

The binarization can be achieved using a manually determined threshold, simple adaptive thresholding methods [169] or more sophisticated techniques such as graph-cuts [1]. However, this approach only yielded reasonable results if the individual objects could still be unambiguously distinguished in the binarized image. To improve the quality of the binarization, to reduce image noise and to possibly emphasize (hyper-) spherical objects of a certain size, a Laplacian-of-Gaussian-based preprocessing was used (Sec. 3.1.2). The possibly preprocessed and binarized image (Fig. 3.6B) was then used to calculate an Euclidean distance map (Fig. 3.6C), where each pixel value of the foreground corresponded to the distance in pixel to the closest background pixel of the binary input image. To guarantee a fast calculation of the EDM, a linear-time algorithm as described by Maurer *et al.* was used [145]. In the ideal case, the distance was maximal in the center regions of the desired objects and the objects were extracted

by subtracting an EDM with suppressed h -maxima from the initial EDM image [210]. The h parameter suppressed local maxima of a height less than h in the input image, *i.e.*, subtracting the EDM with suppressed h -maxima from the initial EDM image yielded only the suppressed h -maxima. The labeled connected components of the binarized h -maxima delivered the actual seed segments (Fig. 3.6D, E). In cases where the objects are more or less equal in size, the h -maxima suppression could be substituted with a binary threshold, which reduced the processing times of the detection. Moreover, if the objects in the original image were well separated, a performance improvement could also be achieved by performing the entire seed detection step on a downscaled version of the image.




Figure 3.6: The most important steps to detect spherical objects in the SBDS1 dataset using an Euclidean distance map-based approach. Based on the raw input image (A), the binarized Laplacian-of-Gaussian of the spherical objects (B) was used to obtain initial object markers. Closely located objects that might be connected due to the binarization could be efficiently split using an Euclidean distance map on the binary images (C). The actual seeds were then extracted from the distance map images by a connected components labeling of the thresholded distance map or the extracted h -maxima of the distance map (D). An overlay of the detected seed segments and the raw image is depicted in (E). This approach is able to separate touching objects as long as the binary region connecting two objects is thinner than the actual object. If this condition is not met, overlapping objects may be detected as a single object.

3.1.4 Validation

Datasets and Compared Algorithms

The performance assessment of the presented seed detection methods was performed on the simulated benchmark datasets SBDS1, SBDS2 and SBDS3 as described in Sec. 3.1.1. In addition to the LoG-based (LoGSM) and the EDM-based (EDM) seed detection methods, another common seed detection

approach based on the connected components of an eroded binary image that has been binarized using Otsu's method (EOCC, [169]) was used for the comparison. Based on the available ground truth, the true positive (TP), false positive (FP) and false negative (FN) detections were counted to calculate recall, precision and the F-Score for each of the methods. True negatives (TN) are in principle all background pixels that were correctly not detected as seed points. However, this information is irrelevant for the seed detection step and was thus omitted. Furthermore, the average Euclidean distance to the correct object centroids as well as the processing time and voxels per second were measured. Additional details on the calculation of the utilized performance measures and the employed testing platform can be found in App. C.1 and App. C.4. Implementation details are provided in Chap. 5.

Performance Evaluation

The seed detection performance of the investigated algorithms obtained on the SBDS1 dataset is listed in Tab. 3.1. All methods provided reasonable results on the benchmark images, where the best values were almost equally shared by the LoGSM and the EDM algorithm. The simpler EOCC approach missed a few objects and additionally produced a small number of false positive detections (7.0% lower F-Score compared to LoGSM and EDM). Regarding the processing times, EDM was the fastest approach, closely followed by the LoGSM method. As the EOCC method was both slower and more inaccurate compared to LoGSM and EDM, the latter methods should be favored. The slower execution time of EOCC was caused by the non-linear morphological erosion operation that cannot be implemented as efficient as a linear LoG filter used in LoGSM or the linear time distance transform performed by EDM.

The influence of different levels of both noise and blur of the images is depicted in Fig. 3.7. EOCC was very sensitive to image noise, as a global binary threshold is always a trade-off between false positive detections and merged objects for noisy images. Hence, the segmentation quality was already affected at the lowest noise levels where a higher amount of merged objects appeared and noise detections increased. At a noise level of $\sigma_{\text{agn}} = 0.004$ no valid segments were found anymore by EOCC. This behavior also explains the improved performance of EOCC for small blur levels, as the Gaussian smoothing serves as a low-pass filter that efficiently suppresses image noise in this case.

Method	TP	FP	FN	Recall	Precision	F-Score	Dist.	Time (s)	KVox./s
EOCC	18.87	1.70	1.13	0.94	0.92	0.93	2.66	1.70	2031.76
LoGSM	19.93	0.00	0.07	1.00	1.00	1.00	1.37	1.34	2577.61
EDM	20.00	0.10	0.00	1.00	1.00	1.00	1.26	1.19	2902.52

Table 3.1: Quantitative performance assessment of the seed points detected in the SBDS1 dataset by an eroded connected components of an Otsu thresholded image (EOCC), a LoG-based multi-scale approach (LoGSM) and an Euclidean distance map-based seed detection method (EDM). Each test image contained exactly 20 objects and the criteria are true positives (TP), false positives (FP), false negatives (FN), recall, precision, F-Score and the distance to the reference (Dist., smaller values are better). Moreover, the achieved average time performance to process a single image was measured in seconds (smaller values are better) and voxels per second (larger values are better). EDM and LoGSM both provided excellent seed detection quality (F-Score = 1) with EDM being faster. As EOCC provided the worst results and was at the same time the slowest method due to non-linear filter operations, it should not be used in practice if the other methods are available.

However, as soon as the variance of the smoothing filter surpassed a value of 10, the detection quality again dropped instantly, as most of the real objects were missed. The variation in the precision values shown in Fig. 3.7E are caused by the adaptive threshold that automatically adjusts its level based on the image content. In contrast to this, LoGSM and EDM both produced reasonable results up to a noise level of $\sigma_{\text{agn}} = 0.01$. Moreover, LoGSM and EDM reliably detected the centroids of the objects up to smoothing variances of 30 and 50, respectively. This is not surprising, however, as the LoGSM method essentially exploits image smoothing to detect the actual object, *i.e.*, an increased blur level up to $\sigma_{\text{smooth}}^2 = 30$ is tolerated by the algorithm. Similarly, the EDM approach worked properly as long as the rough contours of the binarized objects were still distinguishable ($\sigma_{\text{smooth}}^2 \leq 50$), which still resulted in unambiguous intensity maxima of the distance map in the center of the desired objects.

Further validation experiments of the seed detection methods were performed in combination with the uncertainty-based improvements and are covered in Sec. 4.2.1. Additionally, LoGSM and EDM were used to provide seed points for the segmentation approaches discussed in the remainder of this chapter. Thus, the validation of the respective segmentation algorithms also represents an indirect evaluation of the seed detection performance.

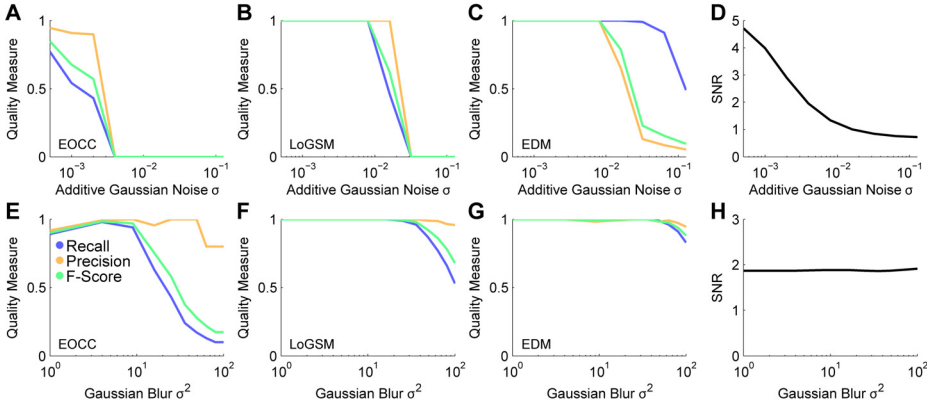


Figure 3.7: Assessment of the seed detection performance of EOCC (A, E), LoGSM (B, F) and EDM (C, G) for different noise levels (A-C) and different Gaussian blur levels (E-G) using the simulated 3D benchmark datasets SBDS2 and SBDS3, respectively. The performance measures recall, precision and F-Score are plotted against the additive Gaussian noise level parameter σ_{agn} (A-C) and the Gaussian blur variance σ_{smooth}^2 (E-G). The influence of the noise and blur level on the signal-to-noise ratio (SNR) of the images is shown in (D, H).

3.1.5 Discussion

In this section, two seed detection methods for the automatic localization of seed points or seed segments in 3D images were presented. Both approaches are well suited to reliably extract the spatial coordinates of desired objects in both noisy and dense environments and the concepts presented for 3D images here can be easily extended to higher or lower dimensions, if necessary. Both approaches were highly robust with respect to increased image noise levels and blurred object contours. On the considered benchmark images, EDM provided the best results. However, due to the single-threaded implementation of the h -maxima filter, the performance of this operator decreased for large images on multicore machines, where the LoGSM operator still provides fast processing speed due to its parallel implementation [220]. The appropriate method has to be chosen according to the field of application. The LoG-based seed detection solely provides the object centroids, which is perfectly suited for object tracking or region of interest determination (Sec. 3.3). Contrary, the EDM-based method provides small seed segments that potentially range over several slices. For instance, for the fusion of 2D segments to 3D objects as performed by the pro-

posed segmentation algorithm for locally plane-like structures (Sec. 3.4). There it was desirable to have seeds spanning over several slices, as this instantly allowed to identify objects that corresponded to a certain segment. All segments that were intersecting with the seed segment could then be merged. Furthermore, having seeds that ranged over multiple slices lowered the chance of missing an object due to a single misdetection of a segment where the seed was located. However, both approaches had difficulties with elongated objects. The LoG-based seed detection tended to extract multiple seeds per desired object at the tips of elongated objects. Furthermore, the elongation occasionally caused intensity plateaus that did not have a single intensity maximum (see Sec. 4.2). Similarly, the EDM-based method tended to detect multiple h -maxima for a single elongated object. The detection behavior could be adjusted by the h parameter of the h -maximum detection that directly depends on the radius of objects of interest by suppressing local maxima of a height smaller or equal to h [210]. Values of h that were smaller than the minimum object radius potentially yielded too many local minima, *i.e.*, split objects. On the other hand, high values of h larger than the minimal object radius tended to merge objects and rejected small objects. If shape and size of the desired objects were comparable, it was usually easy to determine suitable parameter values by manually measuring the radius of the smallest objects that still needed to be detected. However, if object shapes and sizes heavily varied, adjusting this parameter always had to be a trade-off that minimized missing, merged and split objects. In Sec. 4.2.1, an exemplary extension of the LoG-based seed detection method using the uncertainty and fusion framework (Chap. 2) is presented. This extension helped to significantly improve the segmentation results in the case of redundant or missing detections.

3.2 A New Algorithm for the Accurate Extraction and Comparison of Elongated Shapes in 2D Images

Elongated objects are an ubiquitous pattern found in many different image-based experiments. In technical applications, for instance, rods, ropes, cables or pipelines exhibit elongated structures that may need to be compared for au-

omatic quality assurance. In material science, it might be required to quantify cracks in rocks, ice or any other material to draw conclusions about the stability of the investigated substance. Especially, numerous biological tissues such as bones, blood vessels, hair, axons, microtubuli to name but a few, exhibit elongated shapes. Besides extracting the actual elongated objects it is often required to obtain additional information about the local environment, *e.g.*, to quantify dirt particles, neighboring molecules or cells, and basically any other entity that might be present in the close vicinity of the elongated object. In the particular application discussed in this thesis, the analysis was motivated by a toxicological screen of fluorescently labeled neurons in the spinal cord of laterally oriented zebrafish embryos that exhibited both a line-like arrangement of interneurons and blob-like sensory- and motoneurons located above and below the interneurons, respectively (Sec. 6.1, [221]). The specific requirements were not met by any existing software solution and the algorithm was therefore newly developed from scratch. In addition to robustly detecting and extracting the respective objects from the images, a standardized way to compare the extracted information of different images was derived.

3.2.1 Validation Benchmark

To validate the proposed object extraction algorithm, a simulated benchmark that models both an elongated object and randomly distributed small objects in the vicinity of the elongated object was implemented (SBDL). The object types were rendered into two different channels, where channel 1 contained the small objects and channel 2 contained a single elongated object. A random parameterization of a third order polynomial that was sampled over the entire width of the desired output image was used to generate the rough shape of the elongated object. Each ordinate value of the regression curve was further disrupted with a normally distributed random offset. The ground truth image was then formed by iterating over all abscissa values and by adding a vertical bar centered at the ordinate value to form the ground truth image. The ground truth image was subsequently convolved with a Gaussian kernel ($\sigma = 3$) and Poisson noise as well as zero-mean additive Gaussian noise were added to simulate acquisition deficiencies. Finally, an additional channel was added, which contained randomly sized roundish objects that were randomly distributed along

the regression curve. Similar to the elongated object, the small object image was convolved with a Gaussian kernel ($\sigma = 1$) and disrupted using the same noise settings as described before. To generate images with different noise levels, the standard deviation of the zero-mean additive Gaussian noise σ_{agn} was varied in a range $[10^{-5}, 10^{-1}]$. Exemplary images with different levels of additive Gaussian noise are shown in Fig. 3.8 and the benchmark settings are summarized in App., Tab. D.1. Details on the testing platform used for the time performance estimation are provided in App. C.4. For each noise level, 100 different test images were generated and their quality was assessed with respect to the regression curve quality, the correlation between ground truth images and the extracted image regions and the detection quality of the small objects that were randomly distributed along the elongated object.




Figure 3.8: Exemplary images of the SBDL benchmark dataset with different noise levels σ_{agn} . The elongated object is located in the green channel and the spherical objects in the vicinity of the elongated object occupy the other two channels. The standard deviation of the additive Gaussian noise increases from $\sigma_{\text{agn}} = 0$ (A) up to $\sigma_{\text{agn}} = 10^{-2}$ (E). Increased noise levels cause a successively decreasing signal-to-noise ratio and may affect the quality of automatic analysis algorithms.

3.2.2 Algorithmic Design

In order to quantify and identify alterations in the spatial appearance of the different object types with respect to an elongated reference object and to detect variations of the elongated objects themselves, a new image analysis pipeline was developed. The basic steps of the algorithm are summarized in Fig. 3.9 and the individual steps are introduced in the following sections.

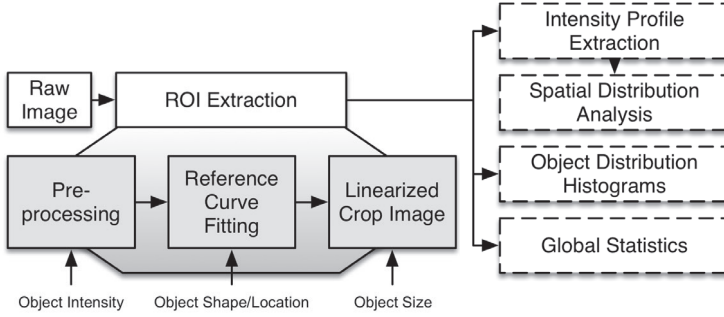


Figure 3.9: Schematic overview of the processing pipeline. The pipeline essentially consists of a regression-based region of interest (ROI) extraction and a straightening of the images to obtain comparable image regions. The actual comparison of extracted crop regions was performed using intensity profiles, spatial distribution analysis, object distribution histograms and general statistical properties of the images. Operators shaded in gray involve prior knowledge about the intensity distribution, object location, the expected size as well as the shape of the object of interest. Result filters are indicated by the dashed edge lines (adapted from [221]).

Region of Interest Extraction

The first step was to extract a comparable image region from the original image I^{raw} with a resolution of $N_x \times N_y$ (Fig. 3.10). If multiple channels were available, this step was performed solely on the channel that contained the elongated object. An initial saturation of the lower q_l -quantile and the upper q_u -quantile of the intensity values was applied to the image for outlier rejection. Furthermore, an additional Gaussian smoothing with standard deviation of σ_{smooth} was used to preserve only the low frequency content. Subsequently, Otsu's method was used to obtain a binarized image from the denoised, smoothed input image (Fig. 3.10B) [169]. A morphological closing was then used to further improve the segmented regions and to obtain a dense connected binary object [210]. If the thresholded image contained multiple connected components, the desired region was selected using prior knowledge about the expected size, location or intensity distribution [221]. If multiple features need to be combined for this regional characterization, the fuzzy set-based uncertainty framework presented in the previous chapter represents a perfectly suited approach for this task. To model the elongated binary object present in the image a polynomial regression curve was fit to the dichotomized image by using the image space coordinates

of the foreground pixels of the object of interest and by minimizing the residual sum of squares. This was essentially accomplished by solving:

$$\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta} \in \mathbb{R}^{N_p}}{\operatorname{argmin}} \sum_{(i_x, i_y) \in \mathcal{P}_{\text{fg}}} (i_y - f(i_x, \boldsymbol{\theta}))^2, \quad (3.4)$$

where f is an arbitrary regression function with unknown N_p -dimensional parameter vector $\boldsymbol{\theta}$ and $\hat{\boldsymbol{\theta}}$ is the optimal parameter vector that minimizes the residual sum of squares. Furthermore, \mathcal{P}_{fg} is the set of foreground pixel coordinates that were used to perform the regression. For the analyses performed in this thesis, a third order polynomial $f(x, \boldsymbol{\theta}) = \theta_1 \cdot x^3 + \theta_2 \cdot x^2 + \theta_3 \cdot x + \theta_4$, with parameter vector $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3, \theta_4)^\top$, was used to sufficiently describe the possible deformations of the imaged specimens. The selection of an appropriate regression model is important to obtain comparable results and has to be matched to the image content. Panels A-C in Fig. 3.10 summarize the individual processing steps performed for the region of interest extraction. The final linearized crop image \mathbf{I}_c^{cr} was formed by extracting a band of predefined radius r below and above the regression curve along the elongated object from the original image. If multiple channels were available, the extraction was performed separately for each of the channels using exactly the same regression curve and extraction radius for all image channels to preserve their colocalization. Fig. 3.10D shows the final result of the ROI extraction on two exemplary benchmark images and has the dimensions $N_x \times (2 \cdot r + 1)$. Global statistical properties of the images were directly extracted from these linearized cropped images. Here, the min, max, mean, median and variance values from the intensity values of each of the channels $c = 1, \dots, N_c$ were extracted for the cropped image to characterize the environment of the elongated object within a distance of r ($\text{gi}_c^{\min}, \text{gi}_c^{\max}, \text{gi}_c^{\text{mean}}, \text{gi}_c^{\text{med}}, \text{gi}_c^{\text{var}}$).




Figure 3.10: Processing steps and exemplary outcome of the elongated shape extraction algorithm applied on two simulated benchmark images. The raw input image containing the elongated object (A, A') was binarized (B, B') and the remaining foreground pixels were used to perform a polynomial regression in a least squares sense (C, C'). A straightened region of interest was extracted along the regression curve using a radius $r = 70$ to obtain a comparable image representation (D, D'). Based on the straightened crop region, intensity profiles (E, E', Sec. 3.2.2) and object distribution histograms (F, F', Sec. 3.2.2) were extracted with respect to the regression curve for further quantitative analysis. The abscissa of both the intensity profiles and the object distribution histograms are depicted relative to the regression curve in pixel (px) with the regression curve being located at the zero mark. Note that the spherical objects are randomly positioned around the zero mark and have a random size, *i.e.*, the intensity profiles and histograms may differ substantially.

Intensity Profile Extraction

To quantitatively assess the distribution of different object types located near or on the elongated object and also to quantify the elongated object itself, the horizontal mean values of each row of pixels in each channel of the cropped image were used to form an average profile vector:

$$\mathbf{p}_c [i_y] = \frac{1}{N_x} \sum_{i_x=1}^{N_x} \mathbf{I}_c^{\text{cr}} (i_x, i_y), \quad (3.5)$$

for all rows $i_y = 1, \dots, 2r + 1$ and all channels $c = 1, \dots, N_c$. This vector \mathbf{p}_c , which is referred to as intensity profile in the remainder of this thesis, can be plotted in a cartesian coordinate system as depicted in Fig. 3.10E.

Another visualization possibility is to plot multiple intensity profiles of different observations vertically aligned to a single heat map image, where each row contains one profile and the intensity levels indicate the height of the respective intensity profiles. This representation was used to easily identify alterations in the intensity distributions of the investigated images [221]. Similar to the global statistical features, single value features like min, max, mean, median and variance of the intensity profiles were extracted ($\text{ip}_c^{\text{min}}, \text{ip}_c^{\text{max}}, \text{ip}_c^{\text{med}}, \text{ip}_c^{\text{mean}}, \text{ip}_c^{\sigma}$). Moreover, calculating the area under the curve yielded a single feature value that quantified the overall signal intensity in the respective channel according to:

$$\text{ip}_c^{\text{auc}} = \sum_{i_y=1}^{(2r+1)} \mathbf{p}_c [i_y]. \quad (3.6)$$

This feature could be refined by using prior knowledge of the spatial location of different objects to constrain the summation only to a desired range. For instance, if a certain class of objects in a specific channel is known to reside above the regression curve, the summation can be adjusted accordingly, in order to use only the relevant information to calculate the respective feature value. In addition to the purely intensity-based features, the characteristics of the intensity profiles, such as peak locations, peak maxima values and the peak width were extracted ($\text{plp}_c, \text{pvp}_c, \text{pwp}_c$).

Spatial Distribution Analysis

The spatial distribution analysis was used to compare the intensity profiles of different experimental groups. In many applications the results obtained for a new experiment need to be compared to a control group $\bar{\mathbf{p}}_c^{\text{ctrl}}$ with known properties. For instance, the control group could be based on manual investigation.

After having defined the control group, all respective intensity profiles of the control group images were combined to a reference intensity profile $\bar{\mathbf{p}}_c^{\text{ctrl}}$. The reference profile was specified by the average of all control profiles with the associated standard deviation series σ_c^{ctrl} as described in [221, 250]. This ensemble averaged reference profile could subsequently be used to identify differences to new observations, *e.g.*, to identify quality deficiencies in automated quality assessment or treatment effects in toxicological experiments. Technically, the comparison was performed by calculating the arithmetic mean of the normalized distance nd_c , which reflects the variance normalized mean deviation of an arbitrary intensity profile \mathbf{p}_c from the averaged reference $\bar{\mathbf{p}}_c^{\text{ctrl}}$:

$$\mu_c^{\text{nd}} = \frac{1}{N_y} \sum_{i_y=1}^{N_y} \text{nd}_c[i_y] = \frac{1}{N_y} \sum_{i_y=1}^{N_y} \frac{|\mathbf{p}_c[i_y] - \bar{\mathbf{p}}_c^{\text{ctrl}}[i_y]|}{\sigma_c^{\text{ctrl}}[i_y]}. \quad (3.7)$$

The feature μ_c^{nd} measures the deviation of a new observation from the averaged reference profile of channel c as n-fold variance. An average value below one indicates that on average all sample points of the investigated intensity profile deviate at most one standard deviation from the averaged reference curve. Hence, a value of zero corresponds to perfect agreement of the compared observation and the reference profile.

Object Distribution Histograms

Another option that yielded more precise information about the spatial distribution of individual objects was the direct extraction of object locations relative to the calculated regression model. This was obtained using blob detection algorithms as described in Sec. 3.1. Extracted object locations were plotted using an object distribution histogram that summarized their frequency with respect to the relative spatial offset from the elongated object. Two exemplary histograms obtained on two validation benchmark images are depicted in Fig. 3.10F. Additional features to quantify the object occurrence were given by the respective mode and maxima values of the histogram. Of course, this approach can be performed independently on multiple colocalized channels if available.

3.2.3 Validation

Datasets and Compared Algorithms

The validation of the proposed segmentation approach was performed on the simulated benchmark dataset SBDL as described in Sec. 3.2.1. The regression-based approach (REG) was parameterized with a standard deviation of $\sigma = 10$, a lower quantile saturation of $q_l = 0.01$, an upper quantile saturation of $q_u = 0.99$ and a region extraction radius $r = 70$. The parameters were empirically determined, such that image noise was suppressed and the object of interest could be unambiguously segmented. The radius was chosen such that the entire region of interest is tightly contained in the cropped image. REG was compared to an alternative approach consisting of a median filter-based noise suppression (5×5 neighborhood), a binary threshold using Otsu's method and an extraction of the elongated object based on the gradient information of the binary image (GRAD).

Performance Evaluation

The quantitative comparison of REG and GRAD is listed in Tab. 3.2 and depicted in Fig. 3.11. Both the elongated structure as well as the surrounding objects were robustly detected by the regression-based approach (REG) up to an additive Gaussian noise level of $\sigma_{\text{agn}} = 10^{-2}$ (SNR: 5 – 6). The gradient-based approach (GRAD) provided reasonable results up to an additive Gaussian noise level of $\sigma_{\text{agn}} = 10^{-3}$ (SNR: 14 – 16). Generally, REG yielded more robust results with lower SSD values of 19 – 21 and higher F-Score values of 0.87 – 0.93 compared to SSD values of 69 – 72 and F-Score values of 0.85 – 0.86 for GRAD. For higher SNR values of 14 – 64, most quality measures remained almost constant. The detection quality of both algorithms, however, abruptly dropped for signal-to-noise ratios lower than 5.9. Starting at an SNR2 value of 5.9, GRAD failed to correctly extract the elongated structures of interest as the kernel size of the involved median low-pass filter could not sufficiently reduce the image noise, whereas REG was still able to correctly identify the region. However, at an SNR2 level of 2.1, REG failed to extract the correct region as well. As both approaches occasionally extracted entirely wrong regions of interest for these low SNR values, the SSD abruptly increased by up to two orders of magnitude and the correlation coefficients dropped to zero.

Method	SNR ₁	SNR ₂	SSD	CC ₁	CC ₂	TP	FP	FN	Rec.	Prec.	F-Sc.	Dist.	t _{roi} (s)	t _{seed} (s)
REG	64	22	19	0.98	0.99	89	1.9	13	0.88	0.98	0.93	1.2	0.39	0.28
GRAD	64	22	69	0.86	0.97	83	8.4	19	0.81	0.91	0.86	1.8	0.31	0.30
REG	41	21	21	0.94	0.98	87	2	12	0.88	0.98	0.93	1.20	0.44	0.28
GRAD	41	21	70	0.79	0.93	80	9	19	0.81	0.90	0.85	1.80	0.31	0.31
REG	16	14	20	0.79	0.87	88	1.7	12	0.88	0.98	0.93	1.20	0.51	0.30
GRAD	16	14	72	0.48	0.69	80	9.40	20	0.80	0.89	0.85	1.90	0.31	0.32
REG	5.10	5.90	21	0.61	0.65	88	13	13	0.87	0.88	0.87	1.30	0.54	0.30
GRAD	5.10	5.90	4817	0.00	0.00	2.40	32	99	0.02	0.00	0.00	-	0.32	0.34
REG	1.80	2.10	1235	0.01	0.02	36	812	65	0.36	0.04	0.08	3.10	0.56	0.31
GRAD	1.80	2.10	4396	0.00	0.00	33	809	68	0.33	0.04	0.07	3.20	0.32	0.35

Table 3.2: Quantitative assessment of the ROI extraction performance of a regression-based approach (REG) and a gradient-based approach (GRAD). Additionally, the object detection quality obtained on the respective linearized crop images is evaluated. SNR1 and SNR2 are the signal-to-noise ratios of channel 1 and channel 2, respectively. SSD reflects the sum of squared distances of the true regression curve and the estimated regression curve. CC1 and CC2 are the correlation coefficients obtained for comparing the extracted crop regions with the ground truth regression curve and the estimated regression curve on the two image channels. The object detection quality was assessed by counting the true positives (TP), the false positives (FP), the false negatives (FN) and by the respective recall, precision and F-Score values. Furthermore, the average distance to the ground truth object centroids was estimated (Dist.). Average processing times for a single image were measured separately for the ROI extraction and the seed detection and are given in seconds.

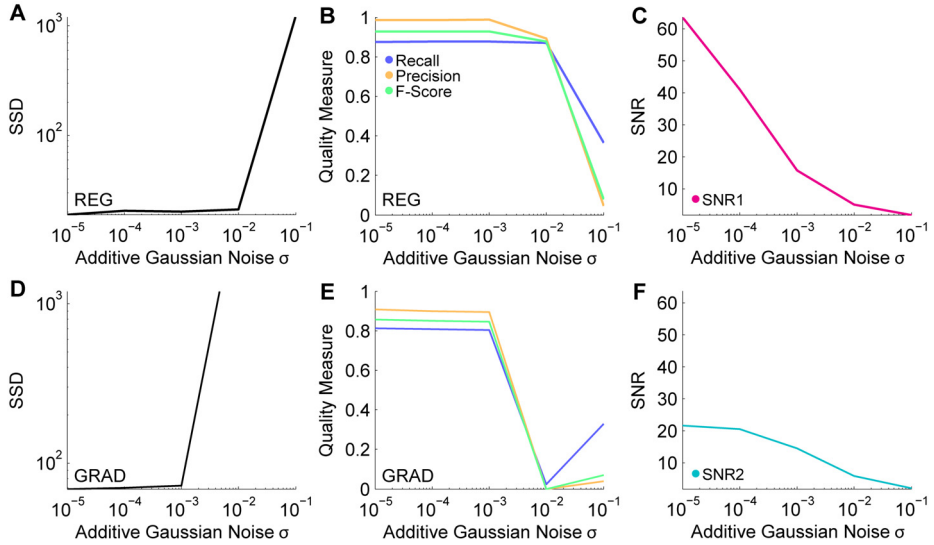


Figure 3.11: Quantitative assessment of the quality of extracted elongated objects as well as the objects in the vicinity of the elongated object using the methods REG and GRAD. The regression curve quality of REG and GRAD was assessed by the sum of squared distances (SSD) of the automatic regression curve and the ground truth regression curve (A, D). The detection quality of objects that were distributed around the regression curve is expressed by recall, precision and F-Score (B, E). The signal-to-noise ratio is plotted in (C, F), where SNR1 and SNR2 were obtained from the small object channel and the elongated object channel, respectively.

The imprecision of the region of interest extraction step directly influenced the spherical object channel as well and caused an abrupt drop of TP values and a heavily increased amount of false positive detections. In some cases, however, the erroneous regions of interest coincidentally overlapped with the true regions of interest. Thus, the quality measures for the seed detection in Fig. 3.11E are not necessarily monotonically decreasing for noise levels $\sigma_{agn} > 10^{-3}$, if parts of the correct region of interest are extracted.

The processing time of REG mainly depends on the number of sample points used for the registration curve fitting, *i.e.*, larger binary regions caused by increased noise levels directly affected the processing time. Contrary, the performance of GRAD is independent of the image content and thus remained constant, even for high noise levels. In both cases, the processing time for the seed detection increased with increased noise levels, as the number of false positive

detections in background regions raised. However, both algorithms required a only a few milliseconds to process a single 2D image, *i.e.*, both of them are well suited even for the analysis of multiple thousands of images, which can be processed in a matter of minutes to maximally a few hours. As REG yielded better results and comparable processing times for the lower noise levels, this method should be favored.

3.2.4 Discussion

In this section, a new approach for the robust extraction of elongated structures in 2D images was presented. Even for images with a low signal-to-noise ratio, a reliable extraction of both the elongated object and the surrounding objects was guaranteed. In comparison to an alternative approach, the regression-based approach was less sensitive to noise and generally extracted more exact regions of interest than the gradient-based approach. Several techniques that were used to compare the extracted image regions between each other have been introduced, namely, intensity profiles, spatial distribution analysis, object distribution histograms and global statistical values of the crop region. Depending on the actual scientific question, the list of extracted features can of course be arbitrarily extended. Although it appears that the object distribution histograms might provide the most quantitative features, the result quality is heavily reduced if the objects of interest are too small or if they form clusters that cannot be reliably resolved with the proposed seed detection methods. However, in these cases, the intensity profiles described in Sec. 3.2.2 essentially provided the same information directly on the intensity level rather than on the object level.

The current version of the algorithm assumed bright objects on a dark background. Depending on the raw image material, additional preprocessing steps such as intensity inversion, contrast adjustment or even more complex methods such as vessel enhancement filters might be necessary to produce reasonable results [64]. The current implementation only searched for a single elongated object of interest in the images. The approach can be adapted to detect and extract multiple elongated structures by simply applying the regression-based ROI extraction on all connected components that contain elongated structures

separately. In cases where the elongated structures overlap, a template matching approach that searches for line-like structures could be employed to perform the initial localization of the objects.

In Sec. 6.1, it is shown how to apply the presented methodology for the extraction of elongated shapes to a toxicological screen to investigate the impact of a large library of small molecules to the neuronal development in the spinal cord of zebrafish embryos.

3.3 A New Algorithm for the Efficient Segmentation of Roundish Objects

A frequently observed segmentation task in many scientific and industrial applications is the detection and segmentation of circular, spherical or generally hyper-spherical objects. Possible objects are, for instance, cell nuclei, stars, nano particles, gas bubbles, coins and basically all other entities that exhibit a (hyper-) spherical shape and are well distinguishable from the background signal. Although the algorithm represents a general methodology to extract sphere-like structures in images of literally arbitrary dimensions, the focus discussed here is put on the fast extraction of fluorescently labeled cell nuclei in 3D microscopy images. Existing methods for this task were not applicable to large-scale time-resolved image data produced by state-of-the-art light-sheet microscopes and it was therefore necessary to find a fast and reliable algorithm to analyze these datasets in a reasonable amount of time [220]. In the remainder of this section the new algorithm is referred to as TWANG (Threshold of Weighted intensity And seed-Normal Gradient dot product image).

3.3.1 Validation Benchmark

The 3D segmentation quality evaluation was performed on the same dataset as described in Sec. 3.1.1, including raw benchmark images (SBDS1) and the disrupted image series with different noise levels (SBDS2) and different blur levels (SBDS3). The specific settings for the respective datasets are summarized in App., Tab. D.1. Additional validation examples performed on a 2D segmentation benchmark are provided in [220].

3.3.2 Algorithmic Design

TWANG is a seed-based segmentation algorithm, which means that the initial locations of the objects about to be extracted should already be roughly known in advance and the respective coordinates have to be provided to the algorithm. In this work, the Laplacian-of-Gaussian scale space maximum projection approach was used for seed detection as described in Sec. 3.1.2. The central idea of the proposed segmentation algorithm is a transformation of the raw input image containing bright spherical objects on a dark background to a representation that can be handled by straightforward adaptive threshold techniques like Otsu's method [169]. A schematic overview of the involved processing steps is shown in Fig. 3.12.

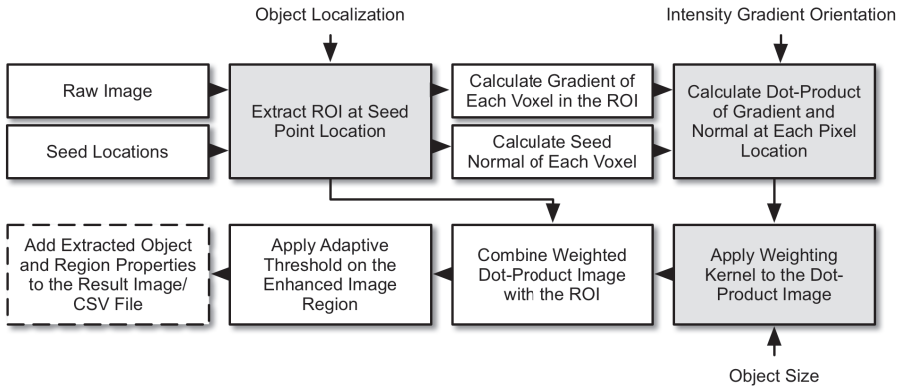


Figure 3.12: Schematic pipeline overview of the TWANG segmentation algorithm. Based on a raw input image and seed point locations of the desired objects, a region of interest of each seed point is transformed to a representation that can be segmented using Otsu's adaptive thresholding technique. All processing steps are performed in parallel for each of the seed points in a data parallel manner. Operators shaded in gray involve prior knowledge such as object localization, intensity gradient orientation and expected object size and the result filter of the pipeline is indicated by the dashed edge.

In order to increase the speed and to be able to detect objects under different illumination conditions, the input image transformation and the segmentation was performed only in the local vicinity of the provided seed points. The first step was to homogeneously distribute the seed points among the available threads and to perform further calculations independently and in parallel. In order to process as few voxels as possible, only small regions surrounding each

seed point were cropped and analyzed. The initial size estimate provided by the seed detection stage was used to estimate the size of this region. For the 3D segmentation, a cuboid with side lengths of $(\sqrt{2} \cdot 3 \cdot \sigma_{\text{seed}}) \cdot (\frac{1}{s_1}, \frac{1}{s_2}, \frac{1}{s_3})^\top$ was used, where $\sqrt{2} \cdot \sigma_{\text{seed}}$ is the radius of the respective seed point at scale σ_{seed} and $\mathbf{s} = (s_1, s_2, s_3)^\top$ corresponds to the physical spacing of the voxels. At each voxel location $\mathbf{x} = (x_1, x_2, x_3)^\top$ of the cropped image \mathbf{I}_s^{cr} of the seed $s = 1, \dots, N_s$, the Gaussian smoothed gradient \mathbf{g}_s was calculated:

$$\mathbf{g}_s(\mathbf{x}, \sigma_{\text{grad}}) = \nabla (\mathbf{I}_s^{\text{cr}}(\mathbf{x}) * \mathbf{G}(\mathbf{x}, \sigma_{\text{grad}})). \quad (3.8)$$

In Eq. (3.8), ∇ is the nabla operator that is used to calculate the partial derivatives of the Gaussian smoothed cropped image. A regularization scale of $\sigma_{\text{grad}} = 3.0$ of the Gaussian kernel \mathbf{G} yielded good results for all investigated scenarios presented in this thesis (see Fig. 3.16 and Tab. 3.3). However, this parameter might need to be adjusted for higher noise levels, in order to guarantee a smooth gradient vector field (see *e.g.*, $\text{SNR} \leq 2$ in Fig. 3.16). By iterating separately over each voxel \mathbf{x} in the sub-region to the respective seed position \mathbf{x}_s , the optimized raw image is sequentially filled with the transformed voxel values. All remaining steps covered in this section are thus defined on the single voxel level. Initially, the difference vector pointing from the seed point location to current voxel was calculated as:

$$\mathbf{d}_s(\mathbf{x}) = \mathbf{s} \circ (\mathbf{x}_s - \mathbf{x}), \quad (3.9)$$

with \circ being the Hadamard product [235], and the seed point normal $\mathbf{n}_s(\mathbf{x})$ at each voxel location is defined as:

$$\mathbf{n}_s(\mathbf{x}) = \frac{\mathbf{d}_s(\mathbf{x})}{\|\mathbf{d}_s(\mathbf{x})\|_2}. \quad (3.10)$$

The seed point normal is a vector pointing from the seed point location to the currently considered voxel \mathbf{x} . Next, the dot product ($\langle \cdot, \cdot \rangle$ -operator) of each normal in the cropped region with the corresponding normalized gradient vector was calculated as:

$$\phi_s(\mathbf{x}) = \frac{1}{2} \cdot \left(1 + \left\langle \frac{\mathbf{g}_s(\mathbf{x})}{\|\mathbf{g}_s(\mathbf{x})\|_2}, \mathbf{n}_s(\mathbf{x}) \right\rangle \right). \quad (3.11)$$

This contrast invariant measure describes the angular dependency of the gradient and the seed point normal and is similar to the one described by Soubies *et al.* [214], where it was used in the energy term of an ellipsoid fit segmentation approach. The transformed dot product of normalized vectors in Eq. (3.11) has a value range of $[0, 1]$, with 1 being identical, 0.5 being orthogonal and 0 being antiparallel vectors. The illustration in Fig. 3.13C, D shows that this property can be exploited to discard voxels in transition regions of touching objects.

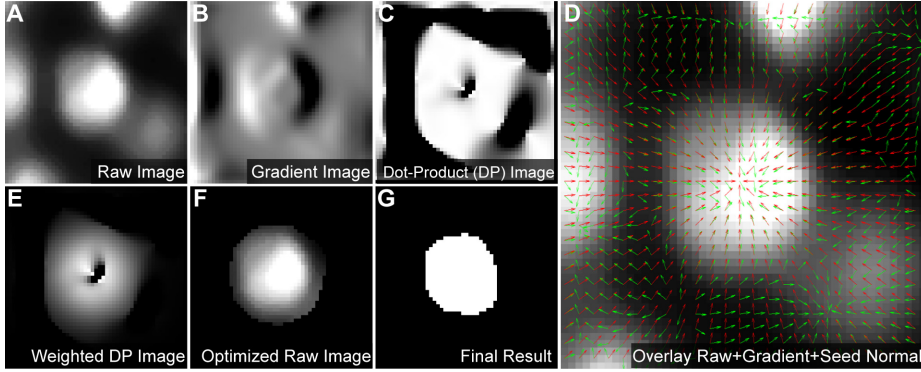


Figure 3.13: Processing steps for the spherical object segmentation on a raw image similar to the SBDS1 dataset. (A) Cropped raw image, (B) Gaussian smoothed left-right derivative image, (C) dot product of the seed normal and the normalized image gradient, (D) overlay of raw image and the smoothed gradient (green) and normal vector (red), (E) weighted dot product image, (F) transformed input image and (G) the final segmentation result (adapted from [220]).

To discard border artifacts and distant neighboring objects, the intensity value of each voxel in the sub-region was decreased if it was far away from the detected seed location. Based on the initial radius approximation $r_s = \sqrt{2} \cdot \sigma_{\text{seed}}$ and a Gaussian kernel standard deviation σ_{kernel} , the following weighting function was employed:

$$w_s(\mathbf{x}) = \max \left\{ \begin{array}{l} \text{sgn} \left((\omega_{\text{kpm}} \cdot r_s) - \|\mathbf{d}_s(\mathbf{x})\|_2 \right) \\ \exp \left(-\frac{((\omega_{\text{kpm}} \cdot r_s) - \|\mathbf{d}_s(\mathbf{x})\|_2)^2}{2 \cdot (\sigma_{\text{kernel}})^2} \right) \end{array} \right. . \quad (3.12)$$

An Gaussian-based smoothing kernel that was calculated using Eq. (3.12) is depicted in Fig. 3.14. The weighting kernel is sampled at the discrete voxel locations of the cropped image, with the weighting kernel's origin being located

in the center of the cropped image. The signum function (sgn) was used to specify the extent of the plateau region in the center of the weighting kernel. The initial size approximation r_s of the seed detection stage (Eq. (3.3)) was used to preserve regions in the close vicinity of the centroid that were smaller than r_s . The weighting kernel plateau radius can optionally be increased or decreased using the multiplier ω_{kpm} with values larger or smaller than the default value of 1.0, respectively. Similarly, the kernel standard deviation σ_{kernel} can be adjusted to flatten the kernel at the borders of the plateau region to increase the influence of voxels at larger distances to the seed center.




Figure 3.14: TWANG segmentation weighting kernel used to discard distant voxels and border artifacts of the cropped region. Exemplary weighting kernels plotted as a line for 1D (A), a surface for 2D (B) and a contour plot for 3D (C), respectively, using the parameters $r_s = 4$, $\sigma_{\text{kernel}} = 3.0$ and $\omega_{\text{kpm}} = 1.0$. The weights are shown on the ordinate axis (A, B) or indicated by the color-code (C). The selected kernel size should encompass the region of interest, such that high weights are obtained in the close vicinity of the object (adapted from [220]).

The weighted normalized dot product image (Fig. 3.13E) could then be obtained by the voxel-wise multiplication of $w_s(\mathbf{x}) \cdot \phi_s(\mathbf{x})$, *i.e.*, a multiplication of the weighting kernel image (Fig. 3.14C) and the dot product image (Fig. 3.13C). The cropped raw image (Fig. 3.13A) and the weighted dot product image (Fig. 3.13E) were combined by copying the original intensity values within the seed radius and by multiplying all remaining raw intensity values with the weighted dot product image. The result of this operation is shown in Fig. 3.13F. Finally, a simple adaptive threshold such as Otsu’s method was applied on the weighted and cropped original image depicted in Fig. 3.13F to obtain the final segmentation shown in Fig. 3.13G. Segment properties such as volume, mean intensity, bounding volumes and the like were immediately extracted from the cropped image regions in parallel. Compared to other segmentation algorithms, a final labeling step on the whole image level can thus be skipped to

accelerate the pipeline execution. Furthermore, the algorithm directly extracts one object at a time per thread based on the predefined size criteria, *i.e.*, literally no merged segments are produced and a watershed-based post-processing to separate merged objects can also be omitted.

3.3.3 Validation

To assess the quality of the new segmentation method (TWANG), a quantitative comparison to Otsu's method (OTSU) [169], Otsu's method with a watershed-based splitting of merged blobs (OTSUWW) [15, 169], a geodesic active contours method (GAC) [39], a gradient flow tracking segmentation (GFT) [125, 133] and a graph cut-based segmentation (GC) [1] was performed. Representative 3D benchmark images of fluorescently labeled nuclei were used (`SBDS1`, `SBDS2` and `SBDS2`), which were similar to the primary application field of the newly developed algorithm. Nevertheless, all illustrations in the following section should be regarded as exemplary use-cases, without the restriction to only work on this particular set of images. In order to work with other imaging modalities or objects, however, it might be necessary to initially transform the input image into a similar representation, *e.g.*, using image inversion, binary thresholds or color space conversions and then applying the proposed algorithm on the transformed image.

Datasets and Compared Algorithms

Otsu's method was selected to demonstrate that the analysis of the respective benchmark images could not sufficiently be analyzed using this adaptive thresholding technique. All other algorithms represent reasonable approaches for the segmentation of fluorescently labeled nuclei, as described in the respective publications [1, 15, 125, 133]. For OTSU and OTSUWW, an additional Gaussian low-pass filter was used to reduce high frequency noise. The level set function of the GAC pipeline was initialized using the LoG-based seed detection method described in Sec. 3.1.2. All algorithms were implemented in C++ (see [220] for implementation details). Details on the performance evaluation measures and the employed testing platform can be found in App. C.2 and App. C.4, respectively. Processing times were measured in seconds.

Performance Evaluation

The quality evaluation for the simulated 3D benchmark data is listed in Tab. 3.3. 3D volume renderings of the false-colored segmentation results are depicted in Fig. 3.15. Merged regions (Fig. 3.15B), split regions (Fig. 3.15E) and missing objects (Fig. 3.15G) are highlighted by the respective symbols. Except of the plain OTSU method (Fig. 3.15B) and the gradient flow tracking method (Fig. 3.15E), most algorithms produced close to perfect topological segmentation results with heterogeneously distributed optimal values (all F-Score values ≥ 0.95). The watershed-based post-processing (OTSUWW) could resolve the erroneous mergers produced by OTSU (Fig. 3.15C) yielding a 15.2% better recall value of 0.99 compared to 0.84 for the plain OTSU. Due to poor edge information, GAC and GFT produced rather too large segments resulting in 1.57 and 1.13 merged objects, respectively. Furthermore, GFT had an increased amount of added objects compared to the other algorithms, which caused the low precision value of 0.71. In both cases, using a different parameterization resulted in heavy under- or over-segmentation. Although regions may not be captured as accurately as with the watershed-corrected adaptive thresholding, the graph-cut implementation provided very good segmentation results and properly split connected nuclei in most cases (F-Score value of 0.97). TWANG offered a high precision (no merged, split or added nuclei) but missed a few cells, yielding a recall of 0.93 and a combined F-Score of 0.96. With respect to the shape based measurements (RI, JI, NSD and HM), OTSU, OTSUWW, GAC and GC produced the best results. The reduced values observed for TWANG and GFT were caused by the increased amount of missing and added nuclei, respectively. However, it can also be observed that OTSU obtained high values in these shape-based categories but is practically not usable without the watershed-based merge correction OTSUWW.

Required processing times significantly varied between the algorithms. The gradient flow tracking (GFT) was up to one order of magnitude slower than OTSU, OTSUWW and TWANG and did not produce better results that would justify the slower execution. When considering the single core performance, the best compromise of speed vs. quality was achieved by the OTSUWW segmentation method. However, this method does not easily scale as no parallel watershed implementation is available yet.

Method	RI	JI	NSD ($\times 10$)	HM	Split	Merg.	Add.	Miss.	Recall	Precision	F-Score	t (s)	t (s)*
OTSU	97.35	6.00	22.82	5.72	0.00	3.27	0.87	0.00	0.84	0.95	0.89	0.49	0.44
OTSUWW	97.57	6.03	3.80	1.12	0.13	0.00	0.00	0.00	0.99	0.95	0.97	2.57	2.48
GFT	88.06	3.57	6.81	6.25	0.10	1.57	6.53	1.87	0.83	0.71	0.77	15.51	-
GAC	95.06	6.40	7.41	2.52	0.00	1.13	0.00	0.77	0.91	1.00	0.95	5.92	-
GC	97.78	6.37	5.66	1.69	1.34	0.07	0.00	0.00	1.00	0.94	0.97	5.92	-
TWANG	93.82	4.94	6.62	2.41	0.00	0.00	0.00	1.37	0.93	1.00	0.96	3.72	1.08

Table 3.3: Segmentation quality of the compared algorithms on the SBDS1 dataset. For quality assessment the Rand index (RI) in percent, the Jaccard index (JI), the Hausdorff metric (HM) and the normalized sum of distances (NSD, values multiplied by 10) as defined in [50] and App. C.2 were used. In addition, the average number of split, merged, erroneously added or missing objects is given. The table summarizes the arithmetic mean values of 30 processed 3D benchmark images. The average performance of the algorithms to process a single image was measured without using threads and with 8 threads where possible (indicated by *) [220]. For GFT, GAC and GC no parallel version could be tested, *i.e.*, the respective performance evaluations are missing.

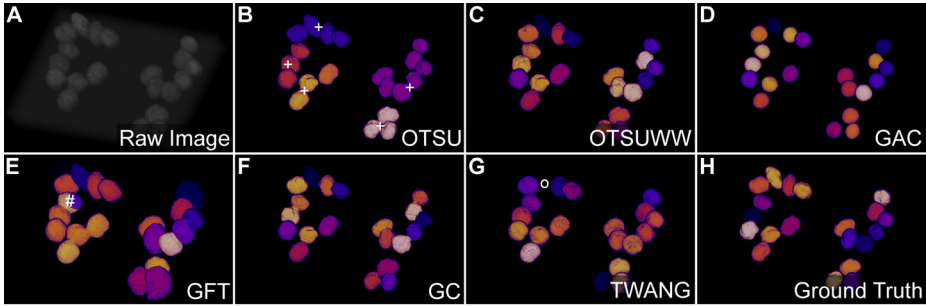


Figure 3.15: Result quality achieved by the compared segmentation algorithms on a set of simulated 3D benchmark images (SBDS1). (A) Simulated original image, (B) adaptive thresholding using Otsu’s method [169], (C) Otsu’s method combined with watershed-based blob splitting [15, 169], (D) geodesic active contours [39], (E) gradient vector flow tracking [125], (F) graph-cuts segmentation [1], (G) TWANG segmentation and (H) the ground truth. Segmentation errors are separated into split (#), merged (+) or missing (o). Each detected object is uniquely colored (adapted from [220]).

Thus, if a multicore processor is available, the TWANG segmentation provides the best compromise between speed and quality.

The three algorithms that provided the best segmentation results on the 3D benchmark (OTSUWW, GC and TWANG) were applied on the disrupted image series to assess their performance on differently flawed images. The recall, precision and F-Score values obtained by the three investigated algorithms on the disrupted image series are shown in Fig. 3.16. All investigated algorithms were at least able to tolerate a small amount of noise ($\sigma_{\text{agn}} \leq 0.001$). The segmentation quality of the GC algorithm, however, already significantly dropped for σ_{agn} levels larger than 0.001. OTSUWW was able to tolerate slightly higher noise levels up to $\sigma_{\text{agn}} \geq 0.002$. Above this level, the global Otsu threshold still recovered the correct segments (recall close to 1 up to $\sigma_{\text{agn}} = 0.01$), but at the same time detected a lot of false positive segments that abruptly lowered the precision of the algorithm. The most robust approach was the TWANG segmentation, which produced acceptable results almost up to noise levels of $\sigma_{\text{agn}} = 0.01$ due to the robust seed detection method used to localize the actual objects of interest. Regarding the increasingly blurred images, OTSUWW and TWANG were able to reliably extract the objects up to a variance $\sigma_{\text{smooth}}^2 = 30$, whereas the quality of GC already slightly dropped at a variance of $\sigma_{\text{smooth}}^2 = 4$.

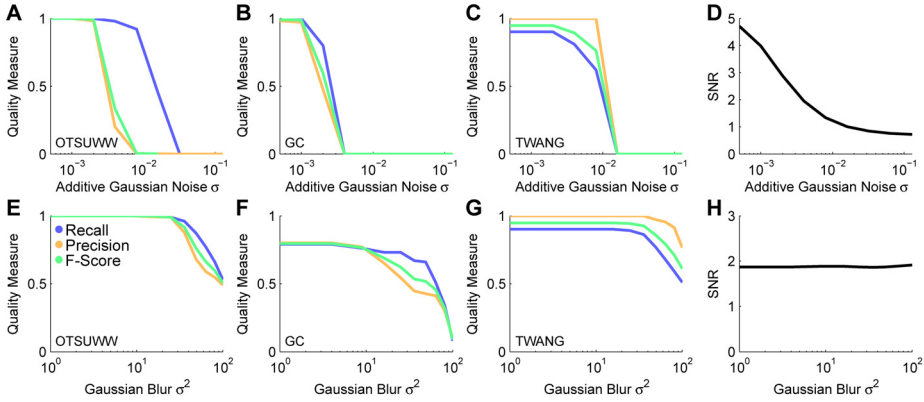


Figure 3.16: Assessment of the 3D segmentation performance of OTSUWW (A, E), GC (B, F) and TWANG (C, G) for different noise and Gaussian blur levels using the simulated 3D benchmark datasets *SBDS2* and *SBDS3*, respectively. The performance measures recall, precision and F-Score are plotted against the additive Gaussian noise level parameter σ_{agn} (A-C) and the Gaussian blur variance σ_{smooth}^2 (E-G). The influence of the noise and blur levels on the signal-to-noise ratio (SNR) of the images is shown in (D, H).

As the blurred images were additionally disrupted with a constant Gaussian noise, the results of GC in Fig. 3.16F already started at a lower level compared to Fig. 3.16B. The results confirm the robustness of the proposed TWANG segmentation and render it a suitable approach even for noisy and possibly blurred images.

3.3.4 Discussion

The presented algorithm for the efficient segmentation of spherical objects in multidimensional images proved to be up to an order of magnitude faster than the compared methods when applied on multicore CPUs. Moreover, the quantitative and qualitative assessment of the segmentation quality showed that the TWANG algorithm produced label images that were comparable to other established methods. Even in the case of severe noise or blurring of the images, the proposed method robustly extracted the desired objects. For the assessment of the segmentation quality on disrupted images, the respective algorithmic parameters were kept constant throughout the experiment. A further extension might be an adaption of the parameter values individually for each

noise or blur level. However, this is almost impossible using manual parameter tweaking and it is necessary to identify (semi-)automated parameter adjustment schemes to accomplish this task.

A remaining problem is the correct extraction of elongated 3D structures that do not match the sphere assumption of the weighting kernel (*e.g.*, if the major axis of the object is twice as large as the two other axes). In these cases, the algorithm tended to clip the tips of ellipsoidal objects due to the spherical shape of the weighting kernel. This behavior could be compensated by adjusting the kernel size multiplier and the kernel standard deviation properly. Further work needs to be put into identifying the elongation properties of the objects already at the seed detection stage, to adapt the weighting kernels accordingly [220].

The quality and reliability of extracted seeds directly influenced the outcome of the segmentation and it was therefore crucial to provide high-quality seeds to the segmentation algorithm. As demonstrated in Sec. 4.2.1, the extended LoG-based seed detection produced excellent results when combined with the uncertainty handling and the fusion of redundant information. The initial size estimate of the multiscale seed detection successfully guided the TWANG algorithm to extract differently sized objects properly [220]. Furthermore, the extracted properties of segmented regions can be used to refine the segmentation results using the uncertainty framework as discussed in Chap. 4. Possible features for the fuzzy set generation are the integrated region intensity, seed intensity, volume, bounding box extents, foreground vs. background ratio and many more.

To sum up, the presented segmentation algorithm is a reasonable choice for a fast analysis of terabyte-scale data that cannot be handled by methods such as [1] due to given memory or time constraints. The presented segmentation algorithm was successfully applied to automatically analyze multiple terabytes of image data of developing zebrafish embryos (Sec. 6.2) in a distributed computing environment based on the ApacheTM Hadoop[®] framework (App. B.2).

3.4 A New Pipeline for Fast and Accurate Segmentation of Locally Plane-like Structures in 3D Image Data

After having introduced algorithms for the seed detection, detection and extraction of elongated objects as well as segmentation of spherical objects, a further frequently observed problem is the segmentation of locally plane-like structures that encapsulate a hollow interior in 3D images. This segmentation task can, for instance, be required to quantitatively assess fluid foams [239], formation of adipose tissue [245] or as in our particular case the segmentation of fluorescently labeled cellular membranes [63, 104, 131, 164]. In Sec. 1.1.2, numerous methods that are in principle suitable for this class of segmentation problems were presented. However, existing methods are either only applicable for 2D images [112] or involve complex processing operators, which in turn significantly limit their application to large-scale microscopy images as produced, *e.g.*, by state-of-the-art light-sheet microscopy [104, 164]. Furthermore, most 3D algorithms require isotropic image data to work properly. As many 3D acquisition techniques produce anisotropic image data, existing approaches perform an upscaling of the images to isotropic resolution, which further increases the processing time required by the algorithms. To overcome these drawbacks of existing algorithms, a hybrid approach of efficient image preprocessing, a slice-based segmentation and a subsequent fusion of extracted 2D segments is proposed, in order to apply the algorithm directly on anisotropic data and to be able to heavily parallelize the performed computations. A schematic overview of the involved processing operators is given in Fig. 3.17 and the new algorithm is referred to as RACE (**R**eal-time **A**utomated **C**ell shape **E**xtractor) in the remainder of this chapter.

3.4.1 Validation Benchmark

A set of simulated images by Mosaliganti *et al.* has been used to evaluate the performance of the developed algorithm on synthetic data of locally plane-like structures in 3D (SBDP1), as well as to be able to assess the segmentation quality under various noise levels and different blur levels [164].

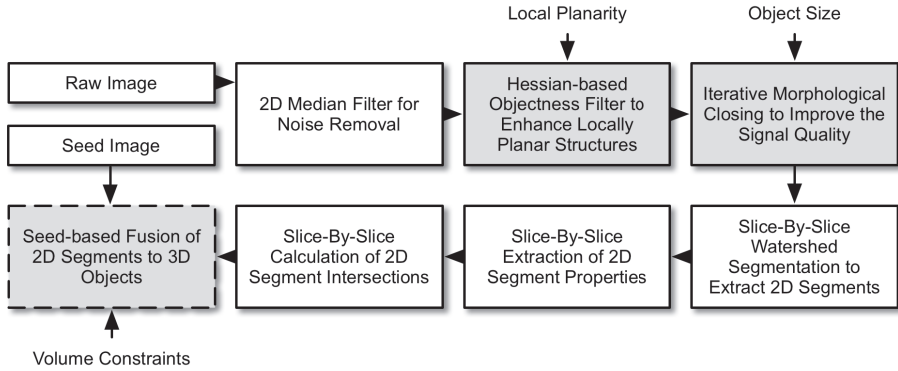


Figure 3.17: Schematic pipeline overview of the RACE algorithm used to efficiently segment locally plane-like structures in 3D images. High-quality 2D segmentations are extracted individually for each of the slices on the basis of an enhanced input image. An additionally provided seed image is used to fuse the 2D segments in each of the slices to form complete 3D objects. Operators shaded in gray involve prior knowledge such as local planarity of desired structures, expected object sizes as well as volume constraints. The output operator is indicated by a dashed edge line.

Similar to the previous sections, the two sets of disrupted images were generated from the simulated benchmark data using zero-mean additive Gaussian noise with $\sigma_{\text{agn}} \in [0.0, 0.128]$ for the first series (SBDP2) and Gaussian blur with variance $\sigma_{\text{smooth}}^2 \in [0, 100]$ and fixed additive Gaussian noise with $\sigma_{\text{agn}} = 0.001$ for the second series (SBDP3). The disruption parameter range was determined manually, to obtain signal-to-noise ratio levels close to the visual detection limit. A set of exemplary benchmark images is shown in Fig. 3.18. The signal-to-noise ratio was determined by the ratio of foreground mean intensity and background intensity standard deviation (Eq. (1.1)). The quantitative measurements were averaged over the quality obtained independently on all annotated images. The segmentation performance was assessed based on the measures described in App. C.2 and the benchmark settings are summarized in App., Tab. D.1. Further validation experiments on real microscopy images are provided in Sec. 6.3 and [216].




Figure 3.18: Exemplary slices of differently flawed 3D benchmark images with different levels of additive Gaussian noise (A-D) and Gaussian smoothed versions with fixed noise level (E-H).

3.4.2 Algorithmic Design

Noise Removal and Hessian-based Ridge Enhancement

As an initial step of the algorithm a shot noise reduction of the input image \mathbf{I}^{raw} was performed using a 2D median filter approach (Fig. 3.19A). For all images considered throughout this thesis, a filter radius of two pixels was sufficient. However, if the images are heavily corrupted by noise, it might be required to use larger filter radii. The median filtered image \mathbf{I}^{med} was further processed using an objectness filter that emphasized locally plane-like structures in 3D images. Here, a generalization of Frangi’s vessel enhancement filter was used [64] as described by Antiga *et al.* [7]. The Hessian matrix was approximated using finite differences of the image intensities of each voxel’s direct neighbors [192]. In order to smooth the response of the filter, the Hessian matrix was calculated on a Gaussian-smoothed version of the input image \mathbf{I}^{med} using a standard deviation σ_{hee} . Based on the eigenvalues of the Hessian matrix of the image at each voxel location, the filter response emphasized M -dimensional objects in N -dimensional images. For the particular case of enhancing locally plane-like structures, *i.e.*, 2D plane-like structures in 3D images, the dimension parameters were set to $M = 2$ and $N = 3$. The generalized objectness function

for this parameterization and the sorted eigenvalues $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$ of the Hessian matrix at each voxel location is given by [7]:

$$O(\boldsymbol{\lambda}(\mathbf{x}, \sigma_{\text{hee}})) = \begin{cases} e^{-\frac{R_B^2}{2\beta_{\text{hee}}^2}} \cdot \left(1 - e^{-\frac{S^2}{2\gamma_{\text{hee}}^2}}\right) & , \lambda_3 < 0 \\ 0 & , \lambda_3 \geq 0. \end{cases} \quad (3.13)$$

In Eq. (3.13), $\boldsymbol{\lambda}(\mathbf{x}, \sigma_{\text{hee}})$ represents the vector of eigenvalues of the Hessian calculated at location \mathbf{x} of the σ_{hee} -regularized input image \mathbf{I}^{med} with the associated Frobenius norm $S = \sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}$ [235]. As a large negative value of the largest eigenvalue λ_3 indicates a sheet-like structure, only this case was considered in Eq. (3.13) [192]. Furthermore, R_B is defined as $|\lambda_2|/|\lambda_3|$ and $\beta_{\text{hee}}, \gamma_{\text{hee}}$ are user-defined weights to control the response of the objectness function (default values proposed in [7]: $\sigma_{\text{hee}} = 10, \beta_{\text{hee}} = 0.5, \gamma_{\text{hee}} = 5.0$). The factor containing R_A , which is found in the original formulation of the generalized objectness filter was omitted here, as it evaluated to one for the considered parameterization [7]. As the default parameters were not applicable for the image material considered in this thesis due to differently sized objects of interest, the parameters were manually optimized to emphasized the structures of interest and set to $\sigma_{\text{hee}} = 2, \beta_{\text{hee}} = 1, \gamma_{\text{hee}} = 0.1$. Sequentially calculating the objectness filter response for all voxels of the median filtered input image yielded an edge-enhanced image \mathbf{I}^{hee} (Fig. 3.19B).




Figure 3.19: Performed processing steps to obtain a high-quality 2D segmentation for each slice of a 3D image stack. Starting with the raw image (A), the locally plane-like signal was enhanced using an eigenvalue-based objectness filter (B). An iterative morphological closing was applied in order to close small gaps in the signal (C). Finally, an accurate 2D segmentation of each slice was obtained using the watershed transform (D). An overlay of the segmentation result and the raw image is depicted in (E) (adapted from [216]).

Iterative Morphological Closing and Slice-based Watershed Segmentation

In order to close small gaps in the locally plane-like structures of the 3D images, an iterative closing of the edge-enhanced image \mathbf{I}^{hee} was performed using grayscale morphology [210]. The morphological closing operation is defined as the erosion of the dilation of the respective input image. In the special case of grayscale morphology, erosion and dilation correspond to the minimum and maximum of the non-zero entries of the structuring element, respectively. Here, an Euclidean sphere was used as structuring element $\mathbf{S}_{r,r}$, where the respective radii r were scaled according to the physical spacing of the input image.

$$\mathbf{I}_r^{\text{cl}} = \begin{cases} \mathbf{I}^{\text{hee}} & , r = 0 \\ \mathbf{I}_{r-1}^{\text{cl}} \bullet \mathbf{S}_r & , r > 0 \end{cases} \quad (3.14)$$

The iteratively closed image up to radius r is denoted by \mathbf{I}_r^{cl} and calculated according to the recursion shown in Eq. (3.14), where \bullet is the morphological closing operation (Fig. 3.19C). The maximum radius of the iterative closing operation was determined by the radius of the smallest objects that should still be detected by the algorithm. Note that the iterative closing can become a bottleneck in the processing pipeline if the used maximum radius is set too large. If only large objects are present in the images it makes sense to down-scale the image instead of using too large radii. Applying a watershed segmentation onto an iteratively closed input image like \mathbf{I}_r^{cl} is also known as the viscous watershed transform as described by Vachier *et al.*, which simulates the flooding of a topographic surface with a viscous fluid and helps to minimize leakage through porous object boundaries [233]. Instead of a physically correct modeling of the viscous fluid, the iterative closing was restricted to a discrete set of predefined radii that was determined by prior knowledge about the desired objects in the images. For all considered images a set of $r = \{1, 2, 3, 4\}$ sufficiently closed gaps in the locally plane-like object signal.

The subsequent watershed transform of the closed image was performed separately on the individual slices, *i.e.*, in 2D instead of directly in 3D as performed by many existing approaches. This yielded a new 3D image \mathbf{I}^{ws} , where each of the slices contained a high-quality 2D segmentation of the desired objects

(Fig. 3.19D, E and Fig. 3.21B). The benefit of the slice-based watershed segmentation was the reduced leakage of background signal into the desired objects, which is a frequently observed issue in 3D watershed segmentations. In addition, the slice-based approach could be perfectly parallelized by performing the calculations of individual slices separately on different threads. The watershed segmentation used in the presented pipeline has a single parameter, which is the intensity dependent starting level of the flooding process, *i.e.*, all local minima below this threshold value are suppressed [15]. Although the 2D segmentation approach helped to speed up the calculations and to mostly get rid of the leakage problem, a subsequent fusion of 2D segments was necessary to obtain the final 3D segmentation as described in the next paragraph.

Seed Point Detection

As the synthetic dataset *SBDP1* only contained simulated membrane structures without central objects, the seeds were directly extracted from this channel. With minor adaptations, this can be achieved by applying the seed detection approach described in Sec. 3.1.3 on an inverted version of the edge-enhanced image I^{hee} to calculate the EDM. In Fig. 3.20, the individual processing steps to extract seeds from the *SBDP1* dataset are illustrated.




Figure 3.20: The most important steps to detect centroids of locally plane-like objects in the *SBDP1* dataset using an Euclidean distance map-based approach. The inverted, binarized version of I^{hee} (B) was used to obtain initial object markers. Closely located objects that might be connected due to gaps in the plane-like objects could be efficiently split using an Euclidean distance map on the binary images (C). The actual seeds were then extracted from the distance map images by a connected components labeling of the thresholded distance map or the extracted h -maxima of the distance map (D). An overlay of the detected seed segments and the raw image is depicted in (E).

Intersection Calculation and Segment Fusion

Based on the slice-based 2D segmentations stored in the I^{ws} image (Fig. 3.21B), the correspondences between slices were determined on the basis of spatial overlap of neighboring segments. Here, the Jaccard index (JI) was used, which is defined as the ratio of the intersection of the two segments vs. the union of the two segments, *i.e.*, it measures the percentage of overlap of two segments with 0 for no overlap and 1 for identical segments. More formally, considering two segments A, B this can be formulated as [247]:

$$JI = \frac{A \cap B}{A \cup B}. \quad (3.15)$$

Assuming that segments with high spatial overlap are likely to belong to the same 3D structure, all intersection pairs were sorted in descending order based on the respective JI similarity value. To avoid fusing very unlikely intersections, a manually determined threshold was used, *i.e.*, only the intersections with values larger than this threshold were considered for further segment fusion. A threshold value of 0.4 produced reasonable results for all segmentation scenarios considered in this thesis. The fusion of the individual segments was obtained by calculating a seeded minimum spanning tree [111]. The seeds used for the segment fusion were derived either directly from the image containing the locally plane-like structures or optionally from an additional image that contained spherical objects that reside inside of the desired locally plane-like objects as described in Sec. 3.1.3. Initially, all 2D segments that intersected with a detected seed were directly labeled with the respective seed label. Intersections within the intersection list where already both segments had a seed label assigned were removed from the list as no further processing was necessary for these segments. All intersections where one of the two segments already had a label assigned were added to an intersection queue, which was again ordered in descending order based on the JI values. The queue was iteratively emptied by popping the top element and by propagating the label of the labeled segment to the unlabeled segment of the intersection. Subsequently, the new segment label was updated for all existing intersections with this newly labeled segment and the respective intersections were inserted into the intersection queue with respect to their similarity value. This process was repeated until the intersec-

tion queue was empty. Initialized by the provided seed points, all 2D segments were enlarged by this technique to the level where no intersections with a similarity above the specified threshold were present anymore (Fig. 3.21C). If the provided seeds perfectly correspond to the number of objects present in the image, no further processing is required at this stage. However, due to potential imperfections of the seed detection, some objects may contain multiple seed points, which results in an over-segmentation of the respective object. To further improve the obtained 3D segmentation results, two fusion heuristics that reduce this kind of over-segmentation errors have been developed. The first heuristic is based on prior knowledge about the expected object size. By specifying a minimum and a maximum expected volume, sub-segments were further fused if they were smaller than the minimum volume and if the fusion additionally did not violate the specified maximum volume value. Similar techniques were also used in existing implementations, however, they were mostly directly performed on the image level rather than on the object level, which can have a significant impact on the time performance of the algorithms.




Figure 3.21: The processing steps to fuse high-quality 2D segments to a complete 3D object. An XZ-slice of an exemplary raw image containing the locally plane-like objects is depicted in (A). Based on the slice-based 2D segmentation (B), a seed-based combinatorial approach was used to merge intersecting 2D segments to single 3D object (C). Two fusion heuristics that incorporated both prior knowledge about intensity and expected object volume were employed to further improve the segmentation results and to correct over-segmentation errors (D). In (E), the final object silhouette is superimposed on the raw input image (adapted from [216]).

The second heuristic is based on calculating the minimum spanning tree directly on the similarity measure according to Kruskal’s algorithm [111] with the additional constraint that no fusion of two segments should introduce a local intensity maximum, which would indicate an object boundary within a segment. The JI-based fusion produced a largely over-segmented result but literally no under-segmentation. Thus, fused segments in the JI-based fusion

were a strong indicator of a required merger of two segments. In cases where a single 3D segment in the JI-based fusion intersected with two segments of the seeded fusion, the two segments of the seeded fusion were combined to a single segment. The final outcome of the fusion algorithm including the proposed fusion heuristics is depicted in Fig. 3.21D.

3.4.3 Validation

Datasets and Compared Algorithms

The proposed RACE segmentation algorithm was compared to two state-of-the-art segmentation algorithms from the field of automatic segmentation of cellular membranes in fluorescence microscopy images. The ACME method applies a tensor voting operation on a ridge-enhanced input image and extracts segments with a 3D watershed algorithm [164]. The MARS method employs alternate sequential filtering combined with a size-dependent iterative watershed transform [63]. The respective version of the RACE algorithm is indicated in round brackets. The specifiers stand for seeds extracted from a spherical objects channel (NS) or seeds obtained directly from the locally plane-like object channel (MS). Moreover, different combinations of enabled and disabled small segment fusion heuristic (SSH) and the Jaccard index-based fusion heuristic (JI) are used. All measurements were performed on the testing platform described in App. C.4. Processing times were measured in seconds and represent the average time needed to process a single image of the dataset.

Performance Evaluation

In Tab. 3.4, the segmentation performance achieved on the SBDP1 benchmark dataset is listed. Moreover, Both ACME and RACE produced almost perfect results with F-Score values of 0.98 and 0.99, respectively. MARS provided a high precision value of 0.99 as it produced almost no split and added cells. On the other hand, due to insufficient border enhancement, MARS merged and missed way more cells compared to ACME and RACE, which resulted in a decreased recall value of only 0.73 and a combined F-Score value of 0.84. The average processing time of RACE was up to an order of magnitude faster compared to the other methods.

Method	RI	JI	NSD ($\times 10$)	HM	Split	Merg.	Add.	Miss.	Rec.	Prec.	F-Sc.	t (s)	s (KVox./s)
RACE _{MS}	99.42	1.00	5.27	1.77	5.25	5.25	1.00	3.00	0.99	0.99	0.99	1.94	430.71
RACE _{MS,SSH}	99.42	1.00	5.27	1.73	2.50	6.00	1.00	3.00	0.98	0.99	0.99	1.69	494.43
RACE _{MS,JI}	99.42	1.00	5.27	1.77	5.25	5.00	1.00	3.00	0.99	0.99	0.99	1.66	503.36
RACE _{MS,JI,SSH}	99.42	1.00	5.27	1.73	2.50	6.00	1.00	3.00	0.98	0.99	0.99	1.71	488.65
ACME	99.48	1.00	2.63	1.39	4.75	10.00	1.00	2.50	0.98	0.99	0.98	28.15	74.21
MARS	98.31	0.98	3.73	1.67	1.25	90.25	1.00	61.25	0.73	0.99	0.84	12.41	168.33

Table 3.4: Quantitative performance assessment of the segmentation methods RACE [216], ACME [164] and MARS [63] on simulated benchmark images [164]. The criteria used to compare the algorithms are the Rand index (RI), the Jaccard index (JI), the normalized sum of distances (NSD) and the Hausdorff metric (HM) as described in App. C.2 and [50]. Additionally, the topological errors were assessed by counting the total number of split, merged, added and missing objects. Precision, recall and F-Score are based on the topological errors by considering split and added nuclei as false positives and merged and missing objects as false negatives, respectively. The average performance to process a single image was measured in seconds (smaller values are better) and voxels per second (larger values are better).

In particular, RACE required less than 2 seconds compared to 28 seconds for ACME and 12 seconds for MARS. The bottleneck of the ACME pipeline is the computationally expensive tensor voting followed by a single-threaded 3D watershed segmentation [164]. The main bottleneck of MARS is as well the single-threaded 3D watershed algorithm for the final segmentation [63]. These costly steps were intentionally omitted in RACE and replaced by a parallel slice-based watershed segmentation. However, the quality benefit of the 3D watershed segmentation and the tensor voting of ACME is reflected in the best scores for the shape-based measurements (RI, JL, NSD and HM). The analysis of the disrupted benchmark image datasets SBDP2 and SBDP3 unveiled that RACE was more robust to image noise, providing excellent results up to $\sigma_{\text{agn}} = 0.002$ and smoothing variances of up to $\sigma_{\text{smooth}}^2 = 4$ (Fig. 3.22A, E). However, for increased blur levels, the quality also significantly dropped. ACME's segmentation quality decreased already at $\sigma_{\text{agn}} = 0.0007$ and was also more sensitive to smoothed object boundaries, which resulted in successively decreased precision already at low smoothing variance values of $\sigma_{\text{smooth}}^2 \geq 1$ (Fig. 3.22B, F). Although the results of MARS were worse compared to ACME and RACE with respect to overall recall value, it was at least able to robustly keep the detection level up to $\sigma_{\text{agn}} = 0.001$ and $\sigma_{\text{smooth}}^2 = 4$, respectively (Fig. 3.22C, G).

In Fig. 3.23, a graphical comparison of the obtained segmentation results is depicted for an exemplary benchmark image. RACE and ACME subjectively produced the best results, which was also confirmed by the quantitative analysis in Tab. 3.4. Both algorithms reliably detected the objects present in the images. Due to the 3D watershed segmentation, the segmentation of ACME qualitatively looks better at the axial transition regions between cells, where RACE missed some slices. However, the general quality was only marginally affected (Tab. 3.4). In contrast to this, MARS was not able to sufficiently close gaps in the membrane signal and produced many merged and erroneous regions. Note that increased blur levels cause an increasing SNR level in Fig. 3.22H, due to a successively decreasing background signal standard deviation compared to an almost constant foreground mean intensity (see Eq. (1.1)). Although there is not yet a clear explanation for the precision variations observed in Fig. 3.22E, this effect was most likely caused by local random noise variations that may have affected the 2D watershed segmentation and in turn the subsequent 3D segment fusion.

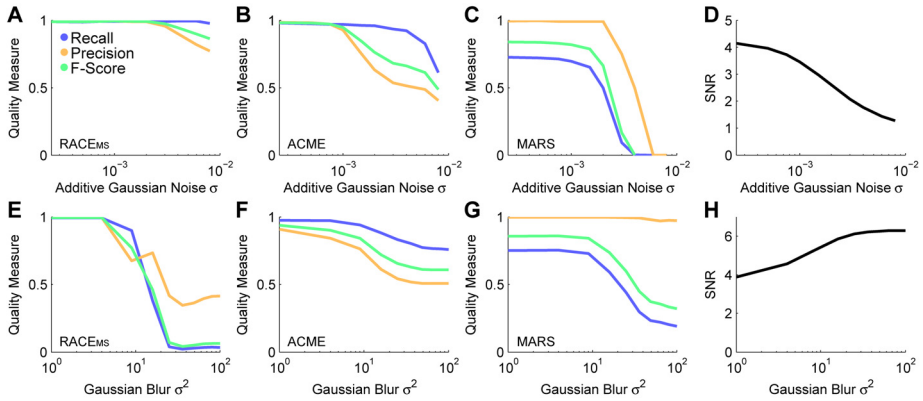


Figure 3.22: Assessment of the 3D segmentation performance of RACE_{MS} (A, E), ACME (B, F) and MARS (C, G) for different noise (SEBP2) and Gaussian blur levels (SEBP3) using simulated 3D benchmark images by Mosaliganti *et al.* [164]. The performance measures recall, precision and F-Score are plotted against the additive Gaussian noise level parameter σ_{agn} (A-C) and the Gaussian blur variance σ_{smooth}^2 (E-G). The influence of the noise and blur levels on the signal-to-noise ratio (SNR) of the images is shown in (D, H).

3.4.4 Discussion

The newly developed RACE algorithm presented in this section proved to provide similar or even better results compared to state-of-the-art algorithms. Even for noisy images or if object contours were slightly smoothed, the algorithm was able to robustly extract the desired objects. At the same time the processing times were at least one order of magnitude lower as for the comparative algorithms. Similar to the TWANG segmentation presented in the previous section, the proposed segmentation algorithm for the reliable extraction of locally plane-like structures was directly dependent on the quality of the provided seeds. It was therefore crucial to tune the seed detection stage to produce as few false positive detections as possible, in order to minimize the amount of split objects in the final segmentation. Possible ways to improve the overall quality of detected seeds in the case of redundant detections was to reject small seeds based on a minimum size criterion or to reject and fuse detected seeds based on prior knowledge-based fuzzy set membership values, as described in Sec. 4.2.

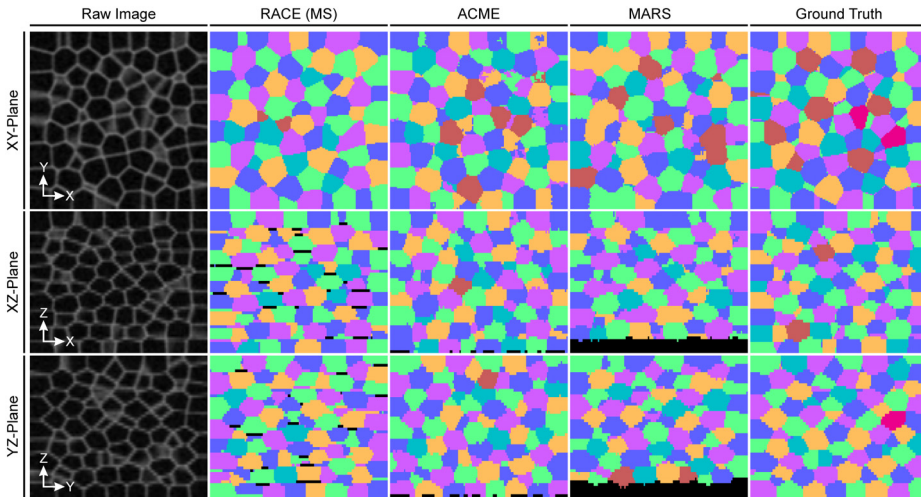


Figure 3.23: Qualitative comparison of the segmentation quality obtained for the different algorithms on one of the *SBDF1* benchmark images. Besides the raw input image, the segmentation results for $RACE_{MS}$, ACME, MARS and the ground truth are shown. The exemplary XY, XZ and YZ slices of the segmentation label images were colored using a graph labeling scheme. Both RACE and ACME provided the best segmentation quality (F-Scores of 0.99 and 0.98, respectively), whereas MARS had issues with many merged objects as well as leakage between the individual objects that lead to many missed cells (F-Score of only 0.84). Due to the 3D watershed segmentation, the axial transitions between cells are more smooth for ACME and MARS compared to RACE (black cavities between cells in XY-Plane and YZ-Plane). However, this influenced the overall quality only marginally (see Tab. 3.4).

The presented segmentation fusion heuristics represent a second feasible approach to partially get rid of over-segmentation errors. Especially, the SSH heuristic was able to successfully correct segmentation errors to a large extent, if object sizes did not vary too much (*e.g.*, developmental stages of an embryo where all cells have comparable volumes that allow one to set tight volume constraints). If object sizes varied heavily (imprecise volume constraints), the improvements achieved by this heuristic were rather limited and it might make sense to add additional heuristics based on object intensity profiles or additional topological criteria. Although, the JI heuristic also slightly improved the segmentation results, its impact was rather small compared to the SSH heuristic [216]. The segmentation quality of all investigated methods was strongly dependent on the sharpness of object contours as shown in Fig. 3.22E-G, *i.e.*, only

relatively small smoothing of the filigree locally plane-like signal was tolerated. It is therefore crucial to ensure that acquired images are perfectly focused and that light scattering is minimized if possible.

Memory limitations that were present, *e.g.*, in the implementation of ACME, were solved by using single precision floating point operations as well as an explicit calculation of the eigenvalues at each pixel location, without calculating a Hessian image for the entire image first. These simple improvements combined with a dedicated GPGPU implementation instantly reduced the required amount of memory by a factor of 18 and enabled processing times that were two orders of magnitude faster compared to existing methods, and eventually enabled real-time data processing [216]. Considering the significant improvement of required processing times as well as the achieved segmentation quality, the presented algorithm is particularly well suited for the automated analysis of large volumetric image stacks, *e.g.*, produced by live imaging of developing model organisms as demonstrated in Sec. 6.3.

4 Enhancing Algorithms with Uncertainty Treatment

In Chap. 2, the translation of available prior knowledge to a mathematical representation using fuzzy sets was introduced. It was shown, how the fuzzy sets can be used in order to determine the validity of extracted information by means of membership degree to a desired class of objects and potential applications such as object filtering, correction and fusion were outlined. Furthermore, three new segmentation algorithms were presented that are usable for efficient information extraction from large-scale multidimensional images (Chap. 3). The present chapter partly combines the results of these two chapters and illustrates how the uncertainty framework can be incorporated into an existing image analysis pipeline for result quality enhancement. The investigated pipeline was inspired by the requirements of the automated analysis of embryonic development and essentially consisted of detection and segmentation of spherical objects (Sec. 4.2, Sec. 4.3) that needed to be combined from different views (Sec. 4.4) and tracked (Sec. 4.5) in time series of 3D microscopy images. An overview of the considered pipeline is provided in Fig. 4.1.




Figure 4.1: Exemplary image analysis pipeline consisting of seed point detection, segmentation, multiview fusion and tracking.

Initially, a comprehensive, generic and artificial benchmark is introduced that served for the validation of all presented processing operators individually and entirely. Moreover, each of the processing operators used for the automatic extraction of spherical objects from three dimensional videos is presented. It is

systematically indicated, at which steps of the pipeline the uncertainty framework was incorporated and how the operators could benefit from the additional information. At first the usage of uncertainties is discussed internally of processing operators, *e.g.*, to exclude unlikely objects from further calculations. In the case of redundant information produced either by applying the same operator on different overlapping views of a scene or by the use of different detection methods on the same image, it is shown how to fuse the respective information appropriately. Furthermore, it is demonstrated how the uncertain information can be shared and propagated between multiple sequentially arranged processing operators to enhance the overall quality of the automated analysis. The processing operators presented in this chapter are entirely motivated by the automated analysis of biological experiments, which are detailed in Chap. 6 but the general concepts can be analogously applied to similar problems from other scientific fields.

4.1 A New Comprehensive Validation Benchmark

To provide a thorough validation of the proposed methods it was inevitable to make extensive use of benchmarking. Various benchmarks that could be employed for the individual components of image analysis pipelines such as seed detection [70, 195], segmentation [50, 195] or tracking [189, 197] have already been presented. An inherent problem of manually generated benchmark datasets, however, is that they often suffer from biased expert knowledge or contain ambiguous image content that is differently rated by different investigators. Hence, simulated data that incorporated all of the mentioned aspects in one comprehensive testing environment was desirable. The charm of using simulated data was the immediate availability of a reliable ground truth and the literally unrestricted possibilities to adjust parameters like noise levels, sampling rates or light attenuation. To overcome the problem of tedious and flawed manual labeling of microscopy images, it has been shown that biological phenomena can be realistically simulated if enough knowledge of the investigated probes was available [118, 120, 144, 189, 224]. However, none of the described approaches could directly be used for the analyses required to evaluate the performance of an entire pipeline comprised of seed detection,

segmentation, multiview fusion and tracking with a single benchmark. The new benchmark required for a comprehensive analysis had to fulfill the following criteria:

- Adjustable image size for empirical processing time estimations
- Realistic simulation of fluorescence properties of labeled nuclei
- Customizable number of cells, nucleus size, cell cycle duration and experimental duration
- Realistic cell behavior, such as cell migration and mitosis events
- Parameterized nucleus dynamics, *e.g.*, different movement models, neighborhood related movement dynamics, spatial restrictions and randomized speed and direction
- Acquisition deficiencies such as an approximated point spread function, slice dependent illumination variations or uneven background illumination
- Simulated multiview generation including light attenuation along the virtual axial direction
- Simulation of detector and discretization related deficiencies like dark current, photon shot noise and signal amplification noise.

To achieve these requirements, existing components and methods were used where possible and condensed with new approaches to a comprehensive framework that provided the desired flexibility. Although, there may be more complex and more realistic approaches to model biologically correct behavior of the cells, the parameters for the generated images were manually tuned to get close to the target datasets of fluorescently labeled nuclei in light-sheet microscopy images while still providing the ease of use, flexibility and speed. This finally enabled testing a multitude of different imaging scenarios in detail and for an entire processing pipeline with a comprehensive benchmark instead of solely assessing the performance of the individual pipeline components on multiple specialized benchmarks. A schematic illustration of the simulated specimen and an overview of the involved simulation steps of the benchmark is shown in Fig. 4.2.

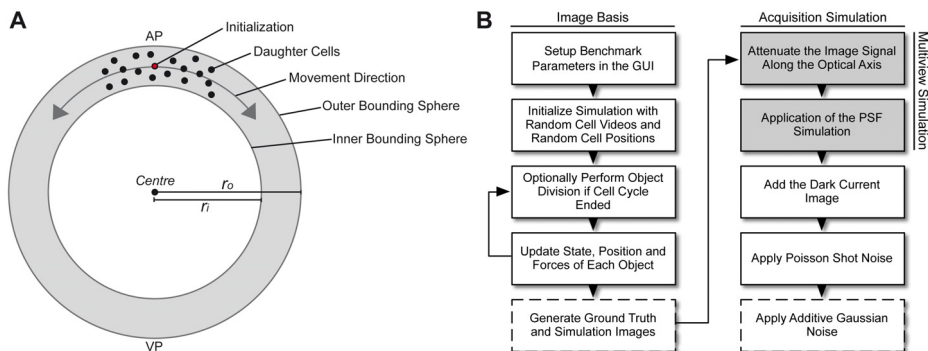


Figure 4.2: (A) Illustration of the embryo simulation. Starting with an initial object, multiple division cycles were simulated including object interaction, object divisions and morphological constraints. Due to inner and outer bounding spheres and the between-object interactions, cells migrated from the animal pole (AP) to the artificial vegetal pole (VP), which produced a behavior similar to the epiboly movement of zebrafish embryo-genesis. (B) The performed steps for a realistic simulated 3D+t benchmark. The left column reflects the object simulation and returned raw images that contained dynamic objects and the associated ground truth data. The right column contains the acquisition simulation, which distorted the simulated images by an artificial signal attenuation, a point spread function simulation (PSF), a dark current image simulation, Poisson distributed photon shot noise and additive Gaussian noise. Steps shaded in gray could optionally consider image rotation, if a multiview experiment was simulated and the output operators are indicated by dashed edge lines (adapted from [219]).

4.1.1 Simulation of Fluorescently Labeled Objects

An object video template database was created from a small simulated time-lapse dataset comprised of eight dividing cells over two division cycles that was kindly provided by D. Svoboda [226]. Sub-videos containing all frames of the division cycles of each cell were extracted and cropped to only contain a single cell (Fig. 4.3A). According to personal communication with D. Svoboda, a higher cell density was not desired for their simulation of suspension cells. Hence, the object video template library described above was used to generate a more challenging benchmark, more precisely, to simulate an artificial developing organism with possibly thousands of dynamic objects. The size of the simulated objects was adjusted to match the size properties observed in representative images of developing zebrafish embryos (Sec. 6.2).




Figure 4.3: (A) Maximum projections of an extracted division cycle of one simulated nucleus [226]. Time increases from left to right and top to bottom. Single objects were randomly initialized and simulated for a predefined experimental duration. This approach yielded a simulated embryo including object movement, object interaction and object divisions with available ground truth (B). Generated raw image sequences were manipulated to simulate various acquisition conditions, such as different levels of additive Gaussian noise (C). The ground truth enabled a quantitative analysis of the algorithmic performance on realistic image data (adapted from [219]).

4.1.2 Generating the Benchmark Images

The first step of the benchmark generation was the setup of the parameters that specified the desired content and the generated ground truth data (Tab. 4.1). Based on the specified image extents a virtual experimental volume was simulated. The selected amount of initial objects was randomly distributed within this experimental volume, each of the created objects got a randomly selected object video assigned and started at a randomly selected cell cycle state. After the initialization, the actual dynamic simulation was started, which was comprised of updating each object's cell cycle state and object video frame. If an object's cell cycle ended during the performed update an object division was performed. Each of the new objects was again randomly initialized and positioned relative to its ancestor. To obtain a dynamically changing scene, the position of each object was updated at every simulation step by considering a set of simulated influences that were acting on it. The interactions were comprised of displacement vectors that originated from adhesive and repulsive behavior between nearby objects, such that objects were not freely floating around but formed clusters.

Simulation Step	Required Parameters	Ground Truth
Image Basis	Initial/Maximum Number of Objects, Maximum Speed of Objects, Movement Model, Time Series Length, Number of Frames, Image Size and Physical Spacing, Object Video Database	Simulated Raw Images, Labeled Mask Images, Object Associations, Extracted Object Properties, Image Statistics
Multiview Simulation	View Transformations, Optical Axis, Simultaneous/Sequential Acquisition	View Transformations, Optical Axis, Simultaneous/Sequential Acquisition
Acquisition Simulation	Light Attenuation Model, Point Spread Function, Dark Frame Intensity, Noise Parameters	Attenuation Model, Point Spread Function, Dark Frame Intensity, Noise Parameters and SNR

Table 4.1: Steps performed for the generation of the simulated benchmark set and the associated ground truth data.

However, objects had to maintain a minimum distance to avoid overlapping of simulated objects. This was accomplished using a repulsive component acting between two objects if their distance became smaller than the sum of their radii. Following the formulation by Macklin *et al.* [139], the displacement vector of the adhesive interaction between two objects was defined as:

$$\Delta \mathbf{x}^{\text{adh}}(\mathbf{d}, R_A) = \begin{cases} \left(1 - \frac{\|\mathbf{d}\|}{R_A}\right)^2 \cdot \frac{\mathbf{d}}{\|\mathbf{d}\|} & 0 \leq \|\mathbf{d}\| < R_A \\ \mathbf{0}, & \text{else.} \end{cases} \quad (4.1)$$

In Eq. (4.1), $\mathbf{d} = \mathbf{x}_j - \mathbf{x}_i$ is the centroid difference vector of two interacting objects i, j and R_A is the maximum adhesive interaction distance [139]. The displacement vector caused by the repulsive interaction was defined as [139]:

$$\Delta \mathbf{x}^{\text{rep}}(\mathbf{d}, R_N, R_M) = \begin{cases} - \left(\left(\left(1 - \frac{R_N}{R_M} \right)^2 - 1 \right) \cdot \frac{\|\mathbf{d}\|}{R_N} + 1 \right) \cdot \frac{\mathbf{d}}{\|\mathbf{d}\|}, & 0 \leq \|\mathbf{d}\| \leq R_N \\ - \left(1 - \frac{\|\mathbf{d}\|}{R_M} \right)^2 \cdot \frac{\mathbf{d}}{\|\mathbf{d}\|}, & R_N < \|\mathbf{d}\| < R_M \\ \mathbf{0}, & \text{else.} \end{cases} \quad (4.2)$$

Next to the centroid difference vector \mathbf{d} , R_N and R_M are the radii of cellular nucleus and the membrane [139]. The influence of the adhesive and the repulsive interaction can be controlled using the weights w_{adh} and w_{rep} , respectively. In addition to the cell-cell interactions, a boundary potential function $\Delta \mathbf{x}^{\text{bdr}}$ was defined to force the objects to move within a narrow volume between two bounding spheres (Fig. 4.2A):

$$\Delta \mathbf{x}^{\text{bdr}}(\mathbf{x}, \mathbf{c}, r_i, r_o, a) = \begin{cases} \frac{\mathbf{x} - \mathbf{c}}{\|\mathbf{x} - \mathbf{c}\|} \cdot \left(1 - \frac{1}{e^{-a(\|\mathbf{x} - \mathbf{c}\| - r_i)}} \right), & \|\mathbf{x} - \mathbf{c}\| < r_i \\ - \frac{\mathbf{x} - \mathbf{c}}{\|\mathbf{x} - \mathbf{c}\|} \cdot \left(1 - \frac{1}{e^{a(\|\mathbf{x} - \mathbf{c}\| - r_o)}} \right), & \|\mathbf{x} - \mathbf{c}\| > r_o \\ \mathbf{0}, & \text{else.} \end{cases} \quad (4.3)$$

In Eq. (4.3), \mathbf{x} is the centroid of the considered object, \mathbf{c} is the center of the bounding volumes, r_i and r_o are the radii of the inner and the outer sphere, respectively, and finally a controls the shape of the sigmoidal boundary potential function. $\Delta \mathbf{x}^{\text{bdr}}$ only contributed to the displacement of a simulated object if the object was already out of the boundary. This additional displacement component prevented objects from entering the inner bounding sphere and from escaping the outer bounding sphere of the simulated embryo (Fig. 4.2A). Similar to the other two interaction types, the influence of the boundary term was controlled using a weight w_{bdr} . In Fig. 4.4, the magnitude of the three displacement terms is qualitatively illustrated.




Figure 4.4: Qualitative magnitude of the boundary displacement (A), the adhesive displacement (B), the repulsive displacement (C) and a combination of the repulsive and the adhesive displacement (D) using the weights $w_{adh} = 0.52$ and $w_{rep} = 1.0$ as proposed in [139]. The boundary displacement vector depends on the distance of the considered object to the center of the bounding sphere. The other two displacement components act between objects, *i.e.*, both depend on the distance between two objects. Note that the repulsive and adhesive displacement vectors point in the opposite direction and the respective influence can be controlled using weight parameters. The intersection of the combined repulsive and adhesive curves with the x-axis shown in (D) represents the equilibrium point of both forces, *i.e.*, objects closer than this distance activate a repulsion and contrary, higher distances trigger an adhesion of two interacting objects.

The total displacement vector of a single object at a given time point can be summarized to:

$$\Delta \mathbf{x}_i^{\text{tot}} = w_{bdr} \cdot \Delta \mathbf{x}^{\text{bdr}}(\mathbf{x}_i) + \sum_{j \in \{1, \dots, N\}}^{i \neq j} \left[w_{rep} \cdot \Delta \mathbf{x}^{\text{rep}}(\|\mathbf{x}_i - \mathbf{x}_j\|) + w_{adh} \cdot \Delta \mathbf{x}^{\text{adh}}(\|\mathbf{x}_i - \mathbf{x}_j\|) \right]. \quad (4.4)$$

Depending on the desired movement model, additional terms can be added to Eq. (4.4), for instance, to add directed movement components or to simulate Brownian motion. However, for the benchmark only passive movements that originated from object density variations, the between-object interaction and the boundary term were considered. As the simulation considered in this thesis was only performed on cell nuclei, the parameters were set in relation to the nucleus radii r_i, r_j of two interacting objects i, j to $R_N = r_i + r_j$ for the nucleus radius parameter and $R_A = R_M = 2 \cdot R_N$ for the membrane and adhesive influence radii. The weights of the displacement components were set to the default values mentioned in [139]: $w_{adh} = 0.52$, $w_{rep} = 1.0$. Moreover, the weight $w_{bdr} = 3.0$ was manually adjusted, such that the interacting objects remained within the spherical boundaries. Note that these parameters

were empirically determined and that the presented model does not necessarily represent an accurate physical simulation of the interacting objects. However, the determined parameters produced movement behaviors that were similar to the epiboly movements observed during early zebrafish development due to increased object densities and boundary constraints that caused a directed movement on a sphere surface. The simulation discussed so far was only performed in the feature space and did not produce any image output. The actual benchmark images and the corresponding label images were initialized as entirely black images. By iterating over all simulated time points and all simulated objects, both the benchmark images and the ground truth images were successively filled with the simulated objects. The respective intensity and label values of the assigned object video frames were simply copied to the result image and the ground truth image, respectively. Note that every object had a unique ID in the label image, to be able to unambiguously identify objects and their frame-to-frame association. A schematic overview of the benchmark generation processing steps is depicted in Fig. 4.2B.

The generation of the benchmark images described so far produced perfectly clear images that were not flawed by any acquisition deficiencies observed in real microscopy images. To approximately model the acquisition process of a real fluorescence microscope, the simulated images were distorted in several steps to obtain a realistic benchmark dataset (Fig. 4.2B). First, the intensities of the simulated images were attenuated with respect to the virtual optical axis. The intensities of each slice were multiplied by a linearly decreasing factor, which was set to 1 at the slice closest to the virtual detection objective and to 0 at the slice farthest from the detection objective. The next step was to convolve the entire image with a point spread function (PSF). Here, the PSFs published in [182] were used, which were measured by imaging fluorescent beads in a light-sheet microscope. Another option frequently found in the literature is to convolve the obtained benchmark images with a three dimensional Gaussian kernel to simulate the point spread function [164]. To generate multiview images with a single rotation of 180° , the multiplier used for signal attenuation was inverted and a point spread function that was analogously rotated by 180° was used to convolve the images.

To simulate the dark current signal of the detector, a positive offset was added to all intensity values. The offset was empirically determined from light-sheet

microscopy data. Furthermore, photon shot noise was simulated by an independent Poisson process at each voxel with the respective image intensities being its average [182]. Finally, the readout noise caused by signal amplification was simulated by zero-mean additive Gaussian noise with a standard deviation of σ_{agn} . Due to the relatively small size of the simulated embryo, light scattering caused by thick tissue was not considered in this benchmark. However, this might be a reasonable extension for upcoming implementations if simulations should get even more realistic.

The individual steps for modeling the acquisition deficiencies of an input image \mathbf{I}^{raw} , which was the raw benchmark image in this case, can be formulated as:

$$\mathbf{I}^{\text{final}} = P_{\lambda}(\mathbf{I}^{\text{raw}} * \mathbf{I}^{\text{psf}} + \mathbf{I}^{\text{dark}}) + \mathcal{N}(0, \sigma_{\text{agn}}), \quad (4.5)$$

where \mathbf{I}^{psf} is the point spread function, \mathbf{I}^{dark} is the dark current image of the detector, P_{λ} applies a Poisson process to simulate photon shot noise and finally $\mathcal{N}(0, \sigma_{\text{agn}})$ is a normally distributed random variable with zero mean and a specified standard deviation σ_{agn} [225].

Based on manually identified parameters for the noise levels and image content, four benchmark datasets were generated for the validation experiments discussed in the following chapters (SBDE1-SBDE4) and the properties of the datasets are listed in App., Tab. D.2. Exemplary benchmark images of different developmental stages and using different signal-to-noise ratios are shown in Fig. 4.3B and Fig. 4.3C, respectively.

4.1.3 Performance Assessment

To assess the performance of the individual image analysis operators, it was necessary to define specific optimality measures. Due to the available ground truth for the generated benchmark datasets it was possible to compare each of the processing steps to the underlying ground truth information. The following measures were used to assess the quality improvements of the seed detection, segmentation, multiview fusion and the tracking. A more detailed description of the employed measures and the testing platform is given in App. C.

Seed Detection

The seed detection quality was evaluated using the benchmark datasets *SBDE1* and *SBDE2* (App., Tab. D.2). The intersections of the detected seeds with the labeled ground truth image were calculated. True positives (TP) were counted as ground truth objects that contained at least one seed point. Seed points that were detected in background regions or redundant detections of ground truth objects were considered as false positives (FP). Ground truth objects that did not contain a seed point were counted as false negatives (FN). Using TP, FP and FN, recall, precision and the F-Score (harmonic mean of precision and recall) were calculated (App. C.1). For all true positives, the average distance to the centroids of the respective ground truth objects was additionally calculated.

Segmentation

The segmentation quality was assessed using the benchmark datasets *SBDE1* and *SBDE2* (App., Tab. D.2). As the provided ground truth of the benchmark contained the complete label images of each frame, a detailed quantitative assessment of the automatic segmentation quality could be performed. The set of segmentation validation measures proposed by Coelho *et al.* was used, namely the Rand index (RI), the Jaccard index (JI), the normalized sum of distances (NSD) and the Hausdorff metric (HM). A detailed description of the measures can be found in [50] and App. C.2. Topological errors produced by the automatic segmentation were separated into added, missing, split or merged objects. Besides the error counts, this topological information was used to define the number of false positives as the sum of split and added cells and analogously the false negatives as the sum of merged and missing cells. These values were then used to calculate recall, precision and F-Score (App. C.1).

Multiview Information Fusion

The quality assessment of information extracted from multiview experiments was based on the *SBDE3* benchmark dataset (App., Tab. D.2). The benchmark dataset contained sequential single view images (SV, one image per time point), sequential multiview images with alternating rotation by 180° (SeMV, one image per time point) and simultaneous multiview images (SiMV, two images per time point). As the simulated objects were moving from frame-to-frame, each

of the complementary rotation images of the SeMV dataset captured a slightly different scene. The evaluation measures were exactly the same as for the segmentation assessment. However, instead of solely comparing a single view to its ground truth dataset, the fused segmentation images of two complementary views were compared to the respective ground truth images. In the case of a sequential acquisition, *i.e.*, an acquisition scheme where the specimen was imaged from one side at a time and was rotated to acquire the opposite views, the results of independently obtained segmentations were fused to a single view. This single fused image was compared to both ground truth frames from different time points individually. In contrast to this, the simultaneous acquisition produced two rotation images at each time point, *i.e.*, here the segmentation results of two views of a single time point were fused and compared to the single ground truth image of this particular time point.

Tracking

To assess the tracking quality, the SBDE4 dataset was used (App., Tab. D.2). The comparison of the investigated algorithms was performed using the TRA measure as described by Maška *et al.* [144] (App. C.3). This measure was calculated by considering the tracking result as an acyclic oriented graph and by comparing this graph to the respective ground truth graph. The inverted, weighted and normalized number of required changes to transform the automatically generated graph to the ground truth graph yielded the normalized TRA measure (higher values are better with 1 being ideal). As the centroids of all objects, the complete temporal association and the object ancestry was known for the simulated SBDE4 dataset, the required ground truth graph could directly be generated using this data. To obtain a more detailed view on the errors made by the respective tracking algorithms, the number of false positive detections, false negative detections, incorrect edges, missing edges, redundant edges and merged objects were counted. Detailed descriptions of the validation measures are provided in App. C.3 and [144].

4.2 Seed Point Detection

In Sec. 3.1.2 a blob detection method that performed an extraction of local extrema in the Laplacian-of-Gaussian maximum intensity projection was introduced. Although, the proposed method worked well in many scenarios, it frequently missed objects that did not exhibit a clear local maximum due to an intensity plateau (*e.g.*, elongated or overexposed objects). To overcome this issue, this section shows how the algorithm was tweaked to be more sensitive to the problematic cases, how redundant seed points could be efficiently fused and how the uncertainty framework proposed in Chap. 2 could be incorporated for efficient data filtering. The introduced steps further improved the automatic detection quality of this algorithm and only had a small impact on the processing times (6% – 35%). Images with varying cell counts (SBDE1) and different noise levels (SBDE2) were used to assess the quality and time performance of the algorithms in different realistic scenarios that were similar to the problems observed in the domain of fluorescence microscopy of living embryos.

4.2.1 Improved Detection and Fusion of Redundant 3D Seed Points

The plain LoG-based seed detection algorithm had difficulties with slightly elongated objects and missed objects that did not have a single strict maximum value, *e.g.*, caused by overexposure of the images. To increase the detection rate even for these problematic cases, the \leq -operator was used instead of the $<$ -operator to identify local extrema, *i.e.*, the detected maxima did not have to be strictly larger than their neighborhood. This approach, however, yielded many false positive detections in background regions and along elongated objects that had to be further filtered and combined in a subsequent processing step.

As an initial step, an intensity threshold (t_{wmi}) was applied to the detected seed points using the window mean intensity feature as described in Sec. 3.1.2. The remaining seed points were mostly located properly on the detected objects and remaining false positive detections largely originated from objects that were detected multiple times. To combine redundant objects to a single one, a fusion approach based on hierarchical clustering was used. The hierarchical cluster tree was computed using Ward’s minimum variance method to compute distances between clusters, *i.e.*, the within-cluster variance was minimized [241].

This method tended to produce clusters that were equally sized and thus rendered it well suited to fuse close maxima that were likely to belong to a single object. The outcome of the hierarchical clustering was a complete tree with an unknown amount of real clusters. Using a distance-based cutoff (t_{dbc}), the final clustering could be obtained from the complete tree. An appropriate cutoff distance could again be derived from prior knowledge about the object size and a good rule of thumb was to set it to the smallest expected object radius r_{min} . Using larger radii can cause a fusion of centroids of neighboring cells and may result in a fused detection in the middle of the centroids of the objects of interest. The feature vectors of all detected seeds within a cluster were then averaged to obtain a single detection per object.

4.2.2 Extending Seed Detection Algorithms by Uncertainty Handling

To improve the performance of down-stream pipeline components such as a multiview fusion operator, the uncertainty of each extracted object at each pipeline stage was additionally estimated as described in Sec. 2.3. The seed detection stage usually represents one of the first analysis steps and therefore, it was assumed that there was no preceding uncertainty information available for the extracted objects. To estimate the uncertainty of the detected seed points, the window mean intensity, the maximum seed intensity and the z-position features of the objects were used. Besides discarding obvious false positive detections in the background regions, the uncertainties were adjusted such that seed detections in low contrast regions (farther away from the detection objective) had lower membership degrees to the class of correct objects than objects in the high contrast regions (closer to the detection objective).

The following steps were performed to obtain the parameterization used for the uncertainty estimation depicted in Fig. 4.5:

- Select the maximum intensity feature and adjust the first two parameters of the trapezoidal membership function to a threshold value that filters all obvious false positives without filtering any true positives.

- Select the seed mean intensity window feature and adjust the first two parameters of the trapezoidal membership function to a threshold value that filters all obvious false positives without filtering any true positives.
- Select the z-position feature and set the first parameter of the trapezoidal membership function to the z-position where the darkest objects that still were detected are present. The second parameter should be set somewhere in the bright-to-dark transition, such that ambiguous detections get lower membership values assigned than obviously correct detections.




Figure 4.5: Visualization of the estimated fuzzy set membership functions for the seed detection stage using the z-position, the maximum seed intensity and the window mean intensity features of each seed point. For the intensity-based features, a fixed threshold was used due to the clear separation of background vs. foreground objects, *i.e.*, the formulation as a fuzzy set was not explicitly required in this case. Seed points are color-coded according to their membership degree to the class of a correct detection and are superimposed on the respective maximum intensity projections along the Z, Y and X axis (A, B, C). The membership functions used for the individual features are depicted in (D) and the obtained membership degrees were multiplied for the final membership degree. The uncertainty gradient along the z-axis was introduced due to the signal attenuation at locations farther away from the detection objective and was used in later steps to resolve ambiguities during multiview fusion (adapted from [219]).

The final fuzzy set membership degree of an object to the fuzzy set of correct objects was determined by the multiplication of all membership degrees of the individual features and was appended as a new feature to the output matrix of the seed detection algorithm (Eq. (2.6)). In this particular example, the seed point maximum intensity and the seed window mean intensity feature are in principal redundant, and could be replaced by only using one or the other. As the window mean intensity should generally be more robust to noise, this feature should be favored. Additionally, no upper limits for the intensity based features were used, which originated from the assumption that the likelihood of a detection being an object of interest does not decrease with increasing intensity. Similarly, the optimal z-position values were not upper-bound as objects closer to the detection objective (higher z-position values in this case) should generally be easier detectable, due to low light attenuation. Of course, these assumptions might have to be adapted if different imaging conditions and specimens are investigated.

The parameterization of the respective fuzzy sets could be easily obtained using the semi-automated analysis approach as described in Sec. 5.2.6. As the GUI provided an instant visual response to parameter changes, it was straightforward to determine the respective transition regions, even for inexperienced users.

4.2.3 Validation

To validate the proposed improvements of the LoG-based seed detection algorithm, the SBDE1 and SBDE2 benchmark datasets were used (App., Tab. D.2). The parameterization and description of the respective versions of the investigated algorithms is listed in Tab. 4.2. Parameters for the scale range to use for the LoG detection were determined empirically, by measuring the radii of the smallest and the largest object that should lead to a valid detection. Here, the minimum and maximum object radii were measured to be $r_{\min} = 5.5$ and $r_{\max} = 11$, respectively. Using the relationship of $\sigma = r/\sqrt{2}$, yielded the standard deviations $\sigma_{\min} = 4$ and $\sigma_{\max} = 8$ for the LoG filtering. The intensity threshold parameters (t_{wmi} , θ_{wmi} , θ_{smi}) as well as the valid z-position range (θ_{zpos}) were identified using the semi-automatic tool described in Sec. 5.2.6.

Method	Parameters	Description
LoGSM	$\sigma_{\min} = 4, \sigma_{\max} = 8, \sigma_{\text{step}} = 1,$ $t_{\text{wmi}} = 0.0025$	Seed detection in the LoG scale space maximum projection with a manually adjusted window mean intensity threshold (t_{wmi}) and a strict maximum detection.
LoGNISM	$\sigma_{\min} = 4, \sigma_{\max} = 8, \sigma_{\text{step}} = 1,$ $t_{\text{wmi}} = 0.0025$	Seed detection in the LoG scale space maximum projection with a manually adjusted window mean intensity threshold (t_{wmi}) and a non-strict maximum detection (Sec. 4.2.1).
LoGNISM+F	$\sigma_{\min} = 4, \sigma_{\max} = 8, \sigma_{\text{step}} = 1,$ $t_{\text{wmi}} = 0.0025, t_{\text{dbc}} = 5$	Same detection as LoGNISM but with additional fusion (F) of redundant detections using a hierarchical clustering approach with a distance-based cutoff value (t_{dbc}).
LoGNISM+F+U	$\sigma_{\min} = 4, \sigma_{\max} = 8, \sigma_{\text{step}} = 1,$ $t_{\text{dbc}} = 5,$ $\theta_{\text{wmi}} = (0.0025, 0.0025, \infty, \infty),$ $\theta_{\text{smi}} = (0.0007, 0.0007, \infty, \infty),$ $\theta_{\text{zpos}} = (50, 250, \infty, \infty),$ $\alpha_{11} = 0.0001, \beta_{11} = \alpha_{11}$	Same detection as LoGNISM but with uncertainty-based (U) threshold and the fusion of LoGNISM+F. The forward threshold α_{11} is set slightly above zero, such that obvious false positives are rejected (Sec. 2.4.1). As no further processing was needed β_{11} was set to α_{11} (Sec. 2.4.2).

Table 4.2: Abbreviations, parameterizations and descriptions of the investigated seed detection algorithms [219].

The extracted seed points of all investigated algorithms were compared to the available ground truth and the quality was assessed using the measures described in Sec. 4.1 and App. C. The obtained values are summarized in Tab. 4.3, whereas each entry of the table corresponds to the arithmetic mean value of the independently obtained results on the ten benchmark images of SBDE1.

Method	TP	FP	FN	Rec.	Prec.	F-Sc.	Dist.	Time (s)	KVox./s
LoGSM	681.1	3.3	202.8	0.77	1.00	0.87	1.60	7.23	7259.82
LoGNSM	813.7	160.0	70.2	0.91	0.84	0.87	1.64	7.66	6858.39
LoGNSM+F	811.7	4.5	72.2	0.91	0.99	0.95	1.59	7.70	6819.58
LoGNSM+F+U	812.5	4.3	71.4	0.91	0.99	0.95	1.59	9.77	6170.05

Table 4.3: Quantitative performance assessment of the LoG-based seed detection methods. The criteria are true positives (TP), false positives (FP), false negatives (FN), recall, precision, F-Score, the distance to the reference (Dist., smaller values are better) as well as the achieved time performance measures in seconds (smaller values are better) and voxels per second (larger values are better). Values represent the arithmetic mean of the results obtained for the individually processed benchmark images (adapted from [219]).

The quantitative analysis confirmed that the proposed extensions of LoGSM could improve the algorithmic performance by up to 9.2% with respect to the F-Score improvement. LoGSM had few false positive detections but on the other hand missed many objects due to the strict maximum detection (recall of 0.77 and precision of 1.0). The recall could be improved by 18.2% to a value of 0.91 by additionally allowing non-strict maxima (LoGNSM). However, this adaptation concurrently raised the number of false positives and thus lowered the precision by 16.0% to 0.84, as objects with maximum plateaus were detected multiple times. These multi-detection errors could be successfully removed using the proposed fusion technique, which was reflected in an F-Score value of 0.95 for LoGNSM+F(+U), *i.e.*, compared to the LoGSM method, the F-Score was increased by 9.2%. Regarding the processing times, the additional effort for a redundant detection was almost negligible, as the non-strict maximum detection simply detected more seed points during the same iteration over the image. The seed point fusion was performed directly in the feature space and was therefore also insignificant compared to the preceding processing steps. For the feature set described here, using the uncertainty-based object rejection (LoGNSM+F+U) only slightly improved the results compared to directly fusing

and filtering the data using the hard intensity threshold but increased the processing time by 35%. LoGNSM+F yielded almost identical results and required only 6% more processing time compared to LoGSM. Nevertheless, all objects were equipped with an uncertainty value that was propagated through the pipeline and proved to be beneficial to filter, fuse and correct the extracted data in subsequent steps, as described in the respective chapters on image segmentation and multiview fusion. In addition, it should be noted that the processing time required for the image analysis easily exceeds the fuzzy set calculations as soon as the images get larger.

To evaluate the influence of the parameter settings on the quality achieved by the different seed detection algorithms, a sensitivity analysis of the most critical parameters for all investigated algorithms was performed on a representative image stack of the benchmark dataset (SBDE1). For the LoG-based seed detection, the intensity-based threshold level was the most critical parameter. Furthermore, the fusion approach additionally required the distance-based cutoff value to be set properly. In Fig. 4.6, the parameter sensitivity analysis of the three investigated LoG-based seed detections is depicted.

The plain LoG with a strict maximum detection yielded high precision values but only a maximal recall of about 0.8 in this case, as objects with intensity plateaus were not detected by the algorithm in contrast to the LoGNSM methods. Contrary to this, the non-strict maximum detection approach almost perfectly detected all objects (recall value of 0.99) but had a lower precision due to multiple detections of a single object that could not be filtered by the simple intensity threshold. Using the hierarchical clustering approach for seed point fusion as described earlier, multiple seed points could efficiently be combined to a single seed point based on a distance parameter. As indicated by the broad parameter range with excellent values of both precision and recall in Fig. 4.6C, the parameter tuning for the seed point fusion is straightforward and the optimal distance can directly be inferred from the minimum expected radius of objects in the investigated specimen (≈ 5 pixel in this example). If the distance for the hierarchical clustering was set too small, no fusions were performed and the results were identical to the LoGNSM method (see decreased quality for distances ≤ 2 in Fig. 4.6C). The detection quality achieved by varying the uncertainty-based threshold is depicted in Fig. 4.6D and yielded the highest quality for membership degree thresholds slightly above zero.

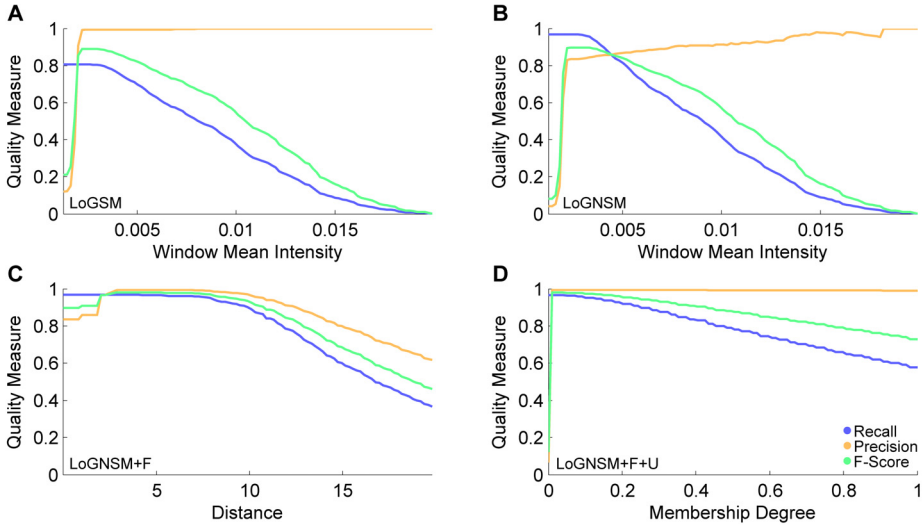


Figure 4.6: Parameter sensitivity analysis for the different versions of the LoG-based seed detection methods. Depicted curves are the achieved recall, precision and F-Score values for different parameter settings. The panels correspond to the plain seed detection with strict maximum detection (A), non-strict maximum detection (B) and a non-strict maximum detection with additional hierarchical clustering for seed point fusion (C, D). In (C) the threshold for the window mean intensity parameter was fixed at 0.0025 and the cutoff distance of the hierarchical clustering was varied. In (D) the threshold was applied on the membership degree with a fixed cutoff distance of 5.

The uncertainty was basically only used to reject highly unlikely objects with an FSMD value of zero ($\alpha_{11} = 0.0001$) and the remaining seeds were propagated unaltered to the next step ($\beta_{11} = \alpha_{11}$). Furthermore, all remaining objects preserved and propagated their FSMD value, which enabled improved decisions at down-stream processing steps.

The performance of the seed detection under different image noise conditions was tested using the SBDE2 dataset, which contained images with different settings for the additive Gaussian noise standard deviation ($\sigma_{\text{agn}} \in [0.0005, 0.01]$). Seed points from these images were extracted using the LoGSM, LoGNSM, LoGNSM+F and LoGNSM+F+U algorithms and the intensity thresholds were determined for each of the noise levels individually using the semi-automatic graphical user interface described in Sec. 5.2.6. An analysis of increasingly blurred images as performed in Chap. 3 was skipped here, as the seed detection

was highly robust up to smoothing levels of $\sigma_{\text{smooth}}^2 = 30$ that are unlikely to be observed in practice (see Fig. 3.7). For higher noise levels, the number of detections in background regions heavily increased and it became ambiguous to determine true positive detections in low contrast regions. The manual threshold was therefore adjusted such that the false positive detections were minimized and only unambiguous seeds were considered. This continuous threshold adaption is also the reason for constant (Fig. 4.7A, C) or increasing precision (Fig. 4.7B), as it was easier to identify false positives rather than false negative detections in noisy image regions. Objects were robustly detected down to a signal-to-noise ratio of 5 (Fig. 4.7), which was close to the visual limit of detection [165] and emphasized the uncertainty-based improvements. The identified seed points and the associated FSMD values were subsequently propagated to the segmentation step to extract a label image of the actual objects as described in the next section.




Figure 4.7: Assessment of the seed detection performance for the different noise levels of the SBDE2 dataset. The performance measures recall, precision and F-Score are plotted versus the additive Gaussian noise level parameter σ_{agn} for LoGSM (A), LoGNMSM (B) and LoGNMSM+F(+U) (C). As LoGNMSM+F and LoGNMSM+F+U produced identical results with respect to recall, precision and F-Score, the plots are combined to a single panel (Tab. 4.3). The influence of the noise level on the signal-to-noise ratio of the images is plotted in (D) (adapted from [219]).

4.3 Segmentation

In Chap. 3, three new algorithms for the automatic segmentation of different geometrical shapes have been presented. The segmentation task required to analyze the simulated benchmark introduced in Sec. 4.1 was to efficiently extract individual spherical objects from simulated 3D image stacks, *e.g.*, using

the algorithms presented in Sec. 3.3. Following the improved seed detection algorithm described in the previous section (LoGNSM+F+U), this section demonstrates how to further improve an algorithm based on adaptive thresholding (OTSUWW). Using propagated information from the seed detection stage and estimated segment FSMDs, both the algorithmic efficiency and the segmentation quality could be improved (OTSUWW+U). Furthermore, the segmentation quality obtained by the TWANG method (Sec. 3.3) was quantitatively compared to a TWANG version (TWANG+U) that relied on the improved seed detection introduced in the previous section (LoGNSM+F+U).

4.3.1 Extending Segmentation Algorithms with the Uncertainty Framework

As an initial step, fuzzy sets for the segmentation evaluation were parameterized in order to estimate the FSMDs of the individual segments to belong to the class of correct objects. Due to the availability of a ground truth, the statistical quantities of the segments could be directly extracted from the reference images. In the absence of ground truth data, the most critical step to manually obtain the fuzzy set parameterization is to estimate the transition regions by a detailed analysis of objects that deviate from the prior-knowledge-based expectation. In the presented case, the focus was put on the volume and size information of the objects in the benchmark images. Thus, to manually determine the described quantities, it is sufficient to investigate a small number of objects at the lower and the upper size and volume level, *e.g.*, using software tools such as Fiji [200]. Of course, the more objects are analyzed, the better the statistics and the resulting fuzzy set membership functions will get. In practice, averaging over about 10 objects at the lower and the upper end, respectively, should be sufficient.

Based on the extracted statistical quantities of the benchmark images that are summarized in Tab. 4.4, the parameter vector $\theta = (a, b, c, d)^\top$ of the trapezoidal fuzzy set membership function for each considered feature was derived using the minimum and maximum values as a, d parameters, respectively. The remaining parameters b, c were set to the 5%-quantile and the 95%-quantile. This parameterization ensured that all values smaller or larger than the maximum values got a membership degree of zero and that 90% of the data range has a

membership value of one assigned. The selection of the transition regions depends on the particular purpose the uncertainty should be used for, *i.e.*, it can be used to tweak algorithms to a certain outcome.

Feature	Min	Max	Mean	Std	Median	5% quantile	95% quantile
Volume	449	2016	993.6	247.2	990	617	1405
Width	13	31	19.9	2.6	20	15	24
Height	13	34	19.9	2.5	20	15	24
Depth	3	11	6.2	1.0	6	5	8

Table 4.4: Statistical quantities of the volume and the extents of all objects contained in the benchmark dataset. Minimum, maximum and quantile values were used to formulate fuzzy sets for each of the features. Individual fuzzy sets were combined using the minimum operator, to obtain a single membership degree value to the fuzzy set of being a valid object (adapted from [219]).




Figure 4.8: Feature histograms (top) and the derived fuzzy set membership functions (bottom) for volume ($\theta_{\text{vol}} = (449, 617, 1405, 2016)^\top$), width ($\theta_{\text{w}} = (13, 15, 24, 31)^\top$), height ($\theta_{\text{h}} = (13, 15, 24, 34)^\top$) and depth ($\theta_{\text{d}} = (3, 5, 8, 11)^\top$). The fuzzy set membership functions were parameterized based on the minimum, maximum values (parameters a, d) as well as the 5% and 95 %-quantile values (parameters b, c) of the objects contained in the benchmark images (adapted from [219]).

As depicted in Fig. 4.8, the shapes of the fuzzy set membership functions derived from the statistical quantities directly resemble the respective distribution observed in the feature histograms. Another option to derive the membership function parameterization is to infer the transition regions via prior knowledge of the investigated specimen. The uncertainty of the output of the automatic

segmentation can directly be assessed by evaluating the membership functions for each feature and by subsequently combining the obtained membership degrees to a single membership value based on the desired combination function (Sec. 2.3.2). Here, the min-operator was used to combine the individual fuzzy set membership degrees to a single value. The combined value directly corresponded to the membership degree of the feature that deviated the most from the specified expected range. Of course, the size criteria discussed here should only be considered as an exemplary illustration. There are various other features that can potentially be used to assess and improve segmentation results, *e.g.*, integrated intensity, edge information, local entropy, local signal-to-noise ratios, principal components, weighted centroids and many more. Furthermore, if colocalized channels are investigated, complementary information can be used to formulate more complex decision rules. In the case of a fluorescently labeled nuclei and membranes that are imaged in different channels (Sec. 3.4), rules like “*each cell has exactly one nucleus*” can be formalized in the same way using the fuzzy set membership functions for a quantification of the available prior knowledge. In the next section the estimated uncertainties will be used to improve imperfect segmentation results and to speed up processing operators.

4.3.2 Uncertainty Guided Segmentation Performance Improvement

The naive application of Otsu’s method on the benchmark dataset produced an immense amount of merged nuclei. Based on the size criteria used to formulate the uncertainty of extracted objects, it was straightforward to determine, which of the objects needed to be further split to match the expected object size. Existing methods usually directly operated in the image domain to improve the segmentation results. For instance, Khan *et al.* [103] and Fernandez *et al.* [63] formulated a minimum volume criterion that was used in a seeded watershed segmentation [15]. Seed points were successively removed from the segmentation, until all segments fulfilled the specified criterion. However, as the methods directly operated in the image domain and applied the costly watershed transformation multiple times, their application to larger 3D images was rather limited [216]. However, if high-quality seed points would be provided to the watershed algorithm, multiple executions could be avoided. This was realized

in the investigated OTSUWW+U implementation using the seed points of the LoGNSM+F+U method described earlier. Additionally, this method could be further improved by applying the watershed-based splitting technique only onto objects that were known to be larger than expected. The estimated FSMD based on the size and volume criteria of manually investigated nuclei could be perfectly exploited to accomplish this task. Using a parallelization strategy similar to the one discussed in Sec. 3.3, all segments with combined FSMD values below β_{21} (backward threshold for operator 2 and linguistic term 1) that corresponded to objects larger than expected (case 3, Sec. 2.3.1) were distributed among the available threads and a seeded watershed approach was used for object splitting in each of the cropped regions of the image. This approach was much faster than directly applying the watershed algorithm on the entire image (OTSUWW), due to the uncertainty guided, locally applied processing of erroneous objects. For simplicity, only a straightforward adaptive thresholding was used here. However, the methods can be analogously applied to erroneous segmentation results produced by more complex algorithms such as level-sets or graph-cuts [1, 123, 218]. After splitting merged objects, the connected components of the image needed to be identified once more and the uncertainty values were re-evaluated as well, to provide the updated information to the subsequent processing operators.

The most critical segmentation error was the presence of many merged regions as described above. However, to further improve the results with respect to false positive detections observed for higher noise levels, segments with combined FSMD values below $\alpha_{21} = 0.1$ (forward threshold for operator 2 and linguistic term 1), *i.e.*, objects that were smaller than the expected object size (case 2 and case 4, Sec. 2.3.1) were removed from the label images. To facilitate the implementation of the uncertainty guided segmentation, it was realized using just a single fuzzy set for the correct class of objects (case 1, Sec. 2.3.1) to identify objects that needed further consideration. To determine if objects with low FSMD values were smaller or larger than the expectation, a comparison to the boundaries of the expected valid range was performed (trapezoidal fuzzy set parameters b, c). Threshold values were again identified using the interactive graphical user interface presented in Sec. 5.2.6.

4.3.3 Validation

To validate the segmentation performance of the described algorithms (OTSU, OTSUWW, OTSUWW+U, TWANG and TWANG+U), all methods were applied to the SBDE1 and the SBDE2 datasets (App., Tab. D.2). In Tab. 4.5, the respective parameterization and a description of each algorithm is provided.

Method	Parameters	Description
OTSU	$\sigma_{\text{smooth}} = 1$	Adaptive thresholding using Otsu's method applied on a regularized raw input image [169].
OTSUWW	$\sigma_{\text{smooth}} = 1$	Same segmentation as OTSU. Merged regions in the entire image were split based on cleaned seeds of the LoGSM method, using a seeded watershed approach.
OTSUWW+U	$\sigma_{\text{smooth}} = 1,$ $\alpha_{21} = 0.1,$ $\beta_{21} = 0.1$	Same segmentation as OTSU. Small noise objects with FSMD values below $\alpha_{21} = 0.1$ were rejected from further processing. Merged regions with FSMD below $\beta_{21} = 0.1$ that were larger than expected were split locally and in parallel using the cleaned and fused seeds of the LoGNSM+F+U method and a seeded watershed approach in small cropped regions.
TWANG	$\sigma_{\text{grad}} = 3,$ $\sigma_{\text{kernel}} = 3,$ $\omega_{\text{kpm}} = 1.41$	TWANG segmentation as described in Sec. 3.3 using the cleaned seeds provided by the LoGSM method. The manually optimized parameter $\omega_{\text{kpm}} = 1.41$ yielded better results than the default value of 1.0 for this dataset.
TWANG+U	$\sigma_{\text{grad}} = 3,$ $\sigma_{\text{kernel}} = 3,$ $\omega_{\text{kpm}} = 1.41,$ $\alpha_{21} = 0,$ $\beta_{21} = 0$	TWANG segmentation as described in Sec. 3.3 using the cleaned and fused seeds provided by the LoGNSM+F+U method. $\alpha_{21} = 0, \beta_{21} = 0$ are adjusted such that all segments are propagated unchanged.

Table 4.5: Abbreviations, parameters and descriptions of the segmentation algorithms [219].

The quantitative segmentation quality results obtained on SBDE1 using the validation measures described in Sec. 4.1 are summarized in Tab. 4.6. The RI value was almost identical for all algorithms and OTSU yielded the highest value. The enhanced adaptive threshold-based techniques yielded an 11.0% better JI value than TWANG+U and the best NSD value was obtained by OTSUWW+U. Considering the results for RI, JI and NSD, the global threshold-based techniques (OTSU*) produced slightly more accurate results (0.3%, 5.3% and 12.0%, respectively) for the objects they were still able to resolve compared to the best results of the TWANG-based methods. Both TWANG-based methods on the other hand produced the minimal amount of topological errors with respect to split and merged objects compared to all OTSU-based methods. The number of added objects was minimal for OTSUWW+U. TWANG+U produced slightly more added objects than TWANG but efficiently detected far more objects. Thus, the F-Score values achieved by TWANG+U were further increased by 8.4% compared to TWANG, *i.e.*, TWANG+U produced the best results with the fewest topological errors (F-Score 0.9). The low amount of split and merged nuclei for TWANG originate from the single-cell extraction strategy, rather than using a global threshold as performed in the OTSU-based methods. The relatively large amount of added objects detected by the TWANG segmentation were mostly no real false positive detections, but segments where most of the extracted region intersected with the image background instead of the actual object and are thus considered as false positives (see low contrast regions in XZ-Projection and YZ-Projection images in 4.9). Regarding the processing times OTSU was the fastest approach, but due to the poor quality without the uncertainty-based extension it was not really an option for a reliable analysis of the image data. TWANG and TWANG+U were 1.8 times slower than the plain OTSU method but 2.4 and 3.8 times faster than OTSUWW and OTSUWW+U, respectively. At the same time, TWANG+U was the most precise approach (F-Score 0.9). Furthermore, OTSUWW+U was 1.6 times faster than OTSUWW due to the focused object splitting and additionally produced better results due to the improved seed detection and noise reduction. These results confirmed that the uncertainty information could be efficiently exploited to guide and constrain computationally demanding processing operators to specific locations and thus, to speed up processing operations while the result quality was preserved or even improved.

Exemplary FSMDs of the final segmentation results of different algorithms are depicted in Fig. 4.9 and provide a convenient visualization to instantly assess the segmentation quality and to identify potential problems of the methods even by non-experts.




Figure 4.9: Maximum intensity projections of the raw image and exemplary volume renderings of the automatic segmentation results produced by OTSU, OTSUWW+U and TWANG+U from different viewpoints (XY, XZ and YZ). The FSMD of individual detected objects was estimated based on morphological criteria such as volume and size (indicated by the color code from red over blue to green for low, medium and high FSMD values) [219].

All algorithms suffered from the light attenuation in the axial direction. Especially, the techniques that relied on a single global intensity threshold had problems to identify the objects located in these low contrast regions. Besides missing many objects, OTSU additionally merged many of the high intensity objects to a single large blob. Especially, in the z-direction many mergers occurred due to the lower sampling in this direction.

Method	RI	JI	NSD ($\times 10$)	HM	Split	Merged	Added	Missing	Rec.	Prec.	F-Score	Time (s)	KVoxel/s
OTSU	97.92	27.82	3.77	6.67	25.00	370.50	42.40	264.00	0.29	0.78	0.42	6.11	8589.46
OTSUWW	97.91	27.99	3.29	6.48	88.60	96.30	57.90	276.20	0.58	0.78	0.67	42.41	1237.72
OTSUWW+U	97.91	27.98	3.15	5.61	50.60	57.80	13.30	268.40	0.64	0.90	0.75	26.40	2035.66
TWANG	97.67	26.59	3.58	4.62	0.00	6.70	75.80	193.50	0.77	0.90	0.83	11.02	4774.83
TWANG+U	97.81	25.30	3.64	4.75	0.00	12.40	103.70	59.50	0.91	0.89	0.90	11.15	4728.84

Table 4.6: Quantitative performance assessment of the segmentation methods OTSU, OTSUWW, OTSUWW+U, TWANG and TWANG+U. The criteria used to compare the algorithms are the Rand index (RI), the Jaccard index (JI), the normalized sum of distances (NSD) and the Hausdorff metric (HM) as described in Sec. 4.1 and App. C. Additionally, the topological errors were assessed by counting split, merged, added and missing objects. Precision, recall and F-Score are based on the topological errors by considering split and added nuclei as false positives and merged and missing objects as false negatives, respectively. The achieved time performance was measured in seconds (smaller values are better) and voxels per second (larger values are better). All values represent the arithmetic mean of the individually processed benchmark images (adapted from [219]).

However, these merged regions could be successfully split to a large extent using the proposed seed-based splitting techniques (OTSUWW, OTSUWW+U). As TWANG directly operated on the provided seeds, it was still able to extract most of the objects in these regions and yielded even higher recall values using the LoGNSM+F+U seed points. However, due to the low contrast, the segmentation quality of the extracted segments in these regions was reduced. In Tab. 4.3, this is reflected by the increased number of added objects for TWANG and TWANG+U, which were mostly no real false positives as described above.

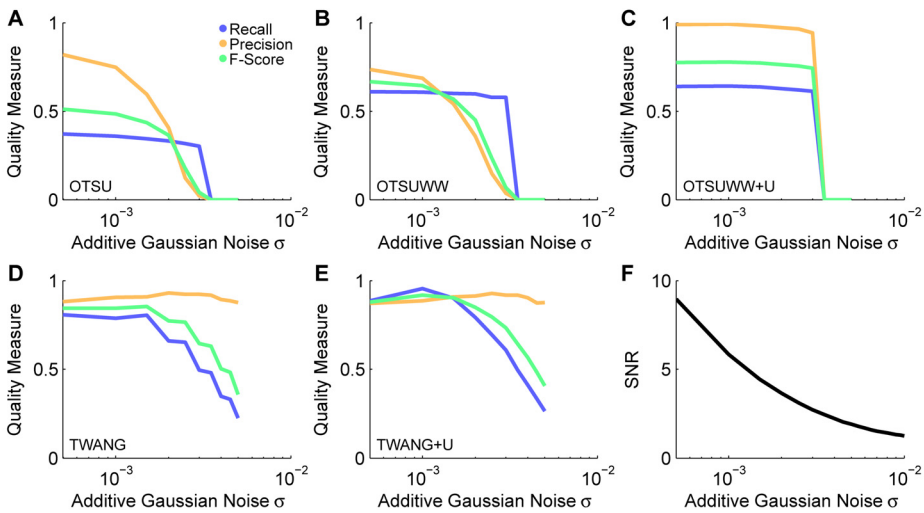


Figure 4.10: Performance evaluation of the segmentation methods OTSU (A), OTSUWW (B), OTSUWW+U (C), TWANG (D) and TWANG+U (E) on images of the SBDE2 dataset with different signal-to-noise ratios. The methods based on adaptive thresholding (OTSU, OTSUWW) suffered from high noise levels and produced a successively increased amount of false positive detections, which could be efficiently suppressed using the uncertainty framework-based extension (OTSUWW+U). The result quality of both TWANG versions directly correlated with the quality of the provided seed points, *i.e.*, TWANG+U benefited from the improved detection rate of LoGNSM+F+U (adapted from [219]).

To investigate the impact of the signal-to-noise ratio of the images to the segmentation quality, the benchmark dataset SBDE2 was processed using all five algorithms. The segmentation quality of all adaptive thresholding-based methods was heavily affected by the noise level of the images yielding poor preci-

sion and recall values even for the lowest noise levels (Fig. 4.10A, B). This was caused by the global threshold that on the one hand merged a lot of objects and on the other hand detected a high amount of false positive segments. The uncertainty-based method OTSUWW+U successfully preserved the increased recall of OTSUWW and at the same time substantially increased the precision to an almost perfect level for noise parameters of $\sigma_{\text{agn}} < 0.003$ (Fig. 4.10C). The increasing number of small segments observed for OTSU and OTSUWW could be efficiently filtered using the uncertainty-based object rejection. As TWANG heavily depended on the quality of the provided seeds, the observed curves in Fig. 4.10C, D show a high correlation to the seed detection performance (Fig. 4.7A, C) and render it as a suitable method even for higher noise levels. The improved seed detection of LoGNSM+F+U also directly affected the quality of the TWANG+U segmentation. Note that the seed points were interactively adjusted for each of the noise levels, *i.e.*, subtle variations in the precision and recall (Fig. 4.10D, E) values are caused by a subjective manual threshold adaption. As in the previous section, increasing blur levels were not considered here, as the seed detection and thus the segmentation worked robustly up to unrealistically high blur levels (see Fig. 3.16). In the following section, the fusion of redundant segmentation information obtained from multiview experiments will be presented. This represents an approach to partially compensate light attenuation and scattering effects that occur when imaging large specimens.

4.4 Extended Multiview Information Fusion

The segmentation results of the approaches presented in the previous section were highly influenced by the decreased image quality along the axial direction. An approach to partly overcome these quality deficiencies caused by attenuation and scattering of light is the acquisition of multiple views from different perspectives. Such complementary image stacks of a specimen can be obtained either by using multiple oppositely arranged detection paths [46, 230] or by a rotation of the probe using a single camera [106]. Both strategies for the acquisition of multiview images require a subsequent fusion step of the information present in the individual views into a single consistent represen-

tation. As a first step, the images have to be aligned properly, *i.e.*, the optimal transformation of the images onto one another has to be identified (Sec. 1.1.2). Most methods proceed with a fusion of the images to form a single high quality image that contains only the best parts of each image and perform all further analysis steps on the fused image [191, 230]. As shown in Sec. 6.2, this fusion of the images may introduce artifacts due to object movements that happen during the sequential acquisition of multiple views. Hence, the approach presented here represents a different way to combine the data obtained from different views, namely the seed point detection and segmentation are applied on the individual rotation images and the multiview fusion of the extracted information is performed after the separately performed image analysis. Naturally, only image regions with sufficient quality should contribute to the fused image and the previously estimated uncertainty information is thus exploited in order to obtain a high quality fusion image.

4.4.1 Uncertainty-based Fusion of Extracted Objects

The fusion of extracted objects can either be purely performed on the feature level or incorporate the segmentation results to identify correspondences between complementary image stacks. The fusion in the feature space can be performed analogously to the fusion of redundant seeds (Sec. 4.2.1) and will thus not be analyzed further in this section. One straightforward improvement of the method presented in Sec. 4.2.1 with respect to multiview acquisition is to fuse neighboring seeds using a convex combination of the detected seeds based on the respective uncertainty values.

Alternatively, a segmentation-based fusion could be performed to obtain a fused segmentation label image. For convenience, it was assumed that the image transformation of 180° was already known in advance and the registration process was therefore skipped. The inputs of the fusion algorithm were two complementary labeled segmentation images of the images that needed to be fused. It was crucial to have unique labels in each of the images to avoid ambiguities of the segment association step. Corresponding segments were identified using a histogram-based approach similar to the one presented in Sec. 3.4.2. However, instead of performing a slice-by-slice intersection calculation, the intersection calculations were directly performed in 3D, *i.e.*, the label

histogram was filled by iterating over the voxels of both images and by successively increasing the 2D histogram bins indicated by the respective image label pairs. Segments that were only present in one or the other image could be found by searching for empty columns and rows of the histogram, respectively, and by copying these segments without any further consideration into the new result image. In the next step, the assignments of the remaining segments in both images were identified by searching for the respective label pair with the maximum overlap. In the case of segmentation fusion it was not desired to perform a weighted average of the segments but rather to use the better segment for the new image. This was accomplished by simply selecting the segment with a higher FSMD value to the desired class of objects. Here, the size-based fuzzy sets specified in Sec. 4.3.1 were used in combination with the information about the axial localization of the extracted objects as specified in Sec. 4.2.2. The uncertainty information was directly provided for each of the extracted segments from the uncertainty quantification performed in the preceding processing steps. Besides the fusion, no further processing of the data was performed at this step, *i.e.*, the uncertainty framework parameters were set to $\alpha_{31} = 0$ (forward threshold of operator 3 and linguistic term 1) and $\beta_{31} = 0$ (backwards threshold of operator 3 and linguistic term 1), to propagate all segments unchanged from the fusion.

4.4.2 Validation

To validate the proposed uncertainty-based multiview segmentation fusion (SF) approach, the two best algorithms identified from the previous section (OTSUWW+U and TWANG+U) were applied to the SBDE3 dataset (App., Tab. D.2). The segmentation quality was assessed for three different acquisition strategies, namely a sequential single view acquisition (SV), a sequential acquisition with a rotation by 180° after each time point (SeMV) and a simultaneous acquisition scheme with oppositely orientated detection paths (SiMV). In the case of SV and SeMV one image was analyzed per time point and two images were analyzed per time point for SiMV. In order to compare the proposed uncertainty-based segmentation fusion approach to a classical intensity-based fusion on the raw image level (IF), an additional set of multi-view images that were combined to a single fused input image using a linear

blending of the intensity values with respect to the rotation of both views was segmented with the OTSUWW+U and the TWANG+U algorithms. In Tab. 4.7, the results of the quantitative analysis are summarized. The SV results were identical to those obtained in the previous section, where especially the global thresholding methods had difficulties to properly extract the segments in low contrast image regions yielding a low recall value of 0.64 for OTSUWW+U. Compared to the SV version of OTSUWW+U, the SeMV version's recall was improved by 32.8% to a value of 0.85 (SF) and by 40.6% to a value of 0.90 (IF), respectively, due to a significantly lower amount of split and missing objects. Although there were slightly more added objects, the split counts also slightly decreased, yielding a precision improvement of 1.1% (SF) and 4.4% (IF), respectively. The result improvements were even more obvious on the SiMV datasets with a greatly improved precision and recall, yielding an increase of the F-Score value by 22.7% (SF) and 25.3% (IF) to 0.92 (SF) and 0.94 (IF) compared to an F-Score of 0.75 obtained by the SV version of the algorithm. Except for the JI value, all other shape dependent measures (RI, NSD, HM) also benefited from the fusion approach. Comparing the results of the SF and the IF fusion approaches, OTSUWW+U worked slightly better on the IF version of the images and obtained the best values for the SiMV datasets. As the segment-based fusion was 3.2 – 5.9-fold slower with the MATLAB-based implementation presented here, the image-based fusion should be preferred if the time resolution is high enough.

The TWANG+U algorithm generally works better in high contrast regions that allow an unambiguous foreground identification. Thus, all shape dependent measures obtained by the TWANG+U segmentation improved in both fusion scenarios, with the best results obtained on the SiMV datasets. Interestingly, the topological segmentation quality of the TWANG+U method was slightly better on the SV dataset as compared to the SeMV dataset (*e.g.*, 1.1% better F-Score on the SV dataset). This was caused by the fact that the seed detection and consequently the segmentation worked already pretty well on the SV images. The comparison of the two different ground truth images to a single fused image therefore performed slightly worse than comparing the two SV segmentations to the two ground truth images with a decreased recall by 2.2% (SF) and a decreased precision by 1.1% (IF).

Method	RI	JI	NSD*	HM	Split	Merg.	Add.	Miss.	Rec.	Prec.	F-Sc.	Time (s)	KVoxel/s
OTSUWW+U (SV)	97.91	27.98	3.15	5.61	50.60	57.80	13.30	268.40	0.64	0.90	0.75	26.40	2035.66
OTSUWW+U+SF (SeMV)	97.93	26.03	3.21	4.91	5.90	43.30	68.80	79.70	0.85	0.91	0.88	66.40/121.40	442.68/809.36
OTSUWW+IF+U (SeMV)	98.05	25.68	3.28	4.62	1.80	64.00	44.00	17.60	0.90	0.94	0.92	21.92	2451.28
OTSUWW+U+SF (SiMV)	98.26	27.19	2.97	4.10	8.90	47.40	4.60	59.10	0.87	0.98	0.92	132.80/242.80	221.34/404.68
OTSUWW+IF+U (SiMV)	98.40	26.97	3.01	3.79	2.50	65.40	3.70	13.70	0.90	0.99	0.94	41.40	1298.10
TWANG+U (SV)	97.81	25.30	3.64	4.75	0.00	12.40	103.70	59.50	0.91	0.89	0.90	11.15	4728.84
TWANG+U+SF (SeMV)	97.89	26.01	3.36	4.64	0.60	5.20	91.00	80.00	0.89	0.89	0.89	51.15/106.15	496.72/1030.82
TWANG+IF+U (SeMV)	97.82	25.47	3.76	4.86	0.40	6.80	105.40	67.20	0.91	0.88	0.89	25.68	2058.21
TWANG+U+SF (SiMV)	98.19	27.06	2.94	3.71	0.30	6.70	21.90	54.00	0.92	0.97	0.94	102.30/212.30	248.36/515.41
TWANG+IF+U (SiMV)	98.14	26.69	3.13	3.83	0.00	7.70	28.10	53.10	0.92	0.97	0.94	26.15	2016.31

Table 4.7: Quantitative performance assessment of the segmentation methods OTSUWW+U and TWANG+U applied on different imaging scenarios. The comparison was performed using sequential single view acquisition (SV), sequential multiview acquisition (SeMV) that rotated the specimen by 180° after each frame and simultaneous multiview acquisition (SiMV) that performed a simulated imaging at each time point with two oppositely arranged detection paths. Images were either fused by the uncertainty-based segment fusion (SF) or using an intensity-based fusion prior to segmentation (IF). The criteria used to compare the algorithms are the Rand index (RI), the Jaccard index (JI), the normalized sum of distances ($NSD \times 10$) and the Hausdorff metric (HM) as described in App. C.2. Additionally, the topological errors were assessed by counting split, merged, added and missing objects. Precision, recall and F-Score are based on the topological errors by considering split and added nuclei as false positives and merged and missing objects as false negatives, respectively. Performing the fusion in the image space took about 15 s per image pair and about 80 – 190 s depending on the number of cells when performed on the segmentation results. The time measurements presented here are the total processing time required to process one frame including segmentation and fusion. Both fusion algorithms were not optimized for speed and implemented in MATLAB. It should be noted that both versions can be easily parallelized and would presumably benefit from a C++-based implementation.

However, this behavior completely changes as soon as the seed detection fails to sufficiently extract the object locations from a single view, where it becomes inevitable to improve the detection quality using multiple views (Sec. 6.2). Applying TWANG+U on the SiMV datasets improved most validation measures in comparison to the SV datasets except of a negligible increased amount of split objects. Especially, the number of added objects could be greatly reduced leading to the best overall F-Score of 0.94. TWANG generally had a lower amount of split and merged cells compared to OTSUWW. As soon as a valid seed point was found, the algorithm successfully extracted the object, due to the explicit prior information about the object size. In contrast to this, OTSUWW initially segmented many merged objects that needed to be split. Some ambiguous split cases that were erroneously split by the watershed algorithm caused the increased number of merged and split objects (20.2% lower F-Score value compared to the best TWANG results). On the other hand, TWANG entirely relies on the seed detection stage and can not directly influence the detection rate and precision. Thus, TWANG tends to have a higher number of added and missing objects compared to OTSUWW (up to 5.3% lower F-Score value in the worst case scenario). The TWANG results obtained using the SF and IF fusion techniques were almost identical, *i.e.*, the uncertainty-based segmentation fusion represents a reasonable alternative to the intensity-based fusion with respect to the segmentation quality. However, with the current MATLAB-based implementation, the time performance of the segment fusion approach was not yet convincing with 3.9 – 8.1-times higher processing times compared to the image-based fusion approach. As the segment fusion approach needed to run two independent segmentations on each of the images, the processing time was already increased 2-fold and required a fusion of the segmented objects afterwards. For both algorithms, approximately 15s were needed for the image-based fusion of two stacks and the segment-based fusion took 80–190s depending on the number of objects. It should be noted, though, that the same implementation in C++ or even using a GPU implementation will certainly make the segment-based fusion competitive and suitable for large-scale analyses.

These results lead to the general conclusion, that a successful fusion of the segmentation results using the proposed uncertainty framework performs reasonably well in all investigated scenarios. As expected, the simultaneous acquisition technique performed better than the sequential acquisition due to the

perfect complementarity of the views. Of course, a sequential multiview acquisition could be improved by a higher sampling rate of the acquired rotation images, such that object movements are compensated sufficiently. Both multiview acquisition approaches significantly improved the segmentation quality obtained by the OTSUWW+U algorithm and finally render this straightforward method with the proposed extensions suitable for the considered image analysis task. Nevertheless, the best choice with respect to a combination of speed and quality was the TWANG+U algorithm that was specifically designed to work on this kind of image data and benefited from both the improved seed points and the additional fusion of complementary frames. The image-based multiview fusion method should be preferred if the temporal resolution is high enough, as the segmentation only needs to be run once in this case. However, if more complex fusion techniques such as wavelet multiview fusion [191] would be required, the segment-based fusion might actually be faster.

A potential extension of the benchmark that is up to future work would be to generate rotation images with different temporal sampling rates, to identify the maximally allowed time shift between the acquisition of sequential multiview data that still yields valid segmentation results. The next section illustrates how fused segmentation results can be used to temporally track dynamic objects.

4.5 Tracking

The final step of the investigated image analysis pipeline was the tracking of all detected objects. As described in Sec. 1.1.2. the basic intention of the tracking procedure is to identify the correct correspondences of detected objects in subsequent time points of the acquired images. For illustration purposes, a straightforward nearest-neighbor tracking approach was used, *i.e.*, for each object present in a frame, the spatially closest object in the subsequent frame was identified and stored as a matching partner. This procedure was applied on every frame of the dataset in order to get a complete linkage of all objects. The tracking algorithm obtained the detected, segmented and fused objects directly from the previous step in a matrix representation similar to Eq. (2.6). The most important features are the centroid of each object as well as the propagated

membership degrees if required for the processing. In addition to the comparison of the segmentation methods introduced in Sec. 4.3, an alternative approach that combines a flawed segmentation with provided seeds is presented in this section (OTSU+NN+U).

4.5.1 Resolving Tracking Errors using Propagated Information

As indicated in the last sections, it was possible to improve segmentation results produced by adaptive thresholding techniques such as OTSU using a watershed-based object splitting (OTSUWW, OTSUWW+U). These improved segmentation results were then propagated to down-stream operators such as multiview fusion or tracking for further processing. Another possibility, however, is to leave the flawed segmentation results produced by the respective algorithms unchanged and additionally provide detected seed points to the tracking algorithm instead of actually using the seed points for splitting in the image domain [218, 219]. For illustration purposes, an Otsu-based threshold was applied to both rotation images independently and the resulting binary images were then fused by simply using the maximum pixel value of the two images. As shown in Tab. 4.6, the segmentation quality achieved by OTSU is in principle not reasonably usable without the watershed-based object splitting. Nevertheless, the tracking algorithm could be extended to decide which of the information was reliable and suitable for tracking and could optionally fall back to the provided seed points if the segmentation quality was insufficient. Here, this was simply achieved by the same fuzzy set classifications as used in chapter Sec. 4.3.1. Using an empirically determined threshold of $\beta_{41} = 0.9$ (backward threshold for operator 4 and linguistic term 1), all objects with an aggregated FSMD value lower than the threshold were not tracked with the actual segment, but with the seed points the respective segment contained. The forward threshold parameter was set to $\alpha_{41} = 0.0$ in order to report all tracking results. Of course, if a further filtering of the obtained tracking results is desired, the algorithms could be adapted appropriately. The next section summarizes the quantitative results obtained by the different pipelines.

4.5.2 Validation

The tracking validation was performed on the SBDE4 dataset, which consisted of 50 frames with two simultaneously acquired rotation images for each frame, yielding a total number of 100 frames that needed to be processed. Segmentation was performed using OTSUWW, OTSUWW+U, TWANG and TWANG+U and all obtained results were fused using the SF approach introduced in Sec. 4.4. The centroids of all detected objects were then used to perform the nearest neighbor tracking (NN). In addition to tracking centroids obtained from the segmentation step, the method described in Sec. 4.5.1 (OTSU+F+NN+U) was applied to the test dataset and the respective tracking results are summarized in Tab. 4.8 and Fig. 4.11.

Both pipelines without uncertainty treatment reached the lowest tracking accuracy with respect to the TRA measure due to an increased number of missing objects (recall of 0.89 for OTSUWW+NN, and 0.88 for TWANG+NN) and a high number of false positive detections (precision of 0.82 for OTSUWW+NN). Of course, these missing objects directly correlated with the number of missing edges and explain the 222.4% (OTSUWW+NN) and 139.9% (TWANG+NN) higher amount of missing edges compared to the best scoring algorithm in this category (OTSU+F+NN+U). Furthermore, OTSUWW+NN suffered from many merged regions that also contributed to the 8.2% lower recall value compared to OTSU+F+NN+U. In contrast to this, all uncertainty-enhanced methods provided comparable results, with the best results achieved by OTSU+F+NN+U and TWANG+U+NN. Although the false negatives of TWANG+U+NN almost halved in comparison to the TWANG+NN method, the detection rate was still the largest problem of this pipeline resulting in a 21.2% higher amount of missing edges compared to OTSU+F+NN+U. On the other hand, OTSU+F+NN+U still suffered from the same under-segmentation tendency as observed for OTSUWW+NN and OTSUWW+U+NN and consequently missing edges were its main problem. With respect to false positive detections both TWANG-based methods provided the best results, with precision values of 1.0 due to a high quality seed detection that benefitted from the simultaneous multiview acquisition. This reflects an improvement of the precision obtained by the TWANG-based methods of 22.0% compared to OTSUWW+NN, which did not have an uncertainty-based object exclusion.

Method	TP	FP	FN	Red.	Miss.	Merg.	Rec.	Prec.	F-Sc.	TRA	Time (s)	KVoxel/s
OTSUWW+NN	905.33	194.10	110.67	5.51	256.88	94.33	0.89	0.82	0.86	0.81	51.36	1020.81
OTSUWW+U+NN	941.94	65.65	74.06	2.63	141.39	55.04	0.93	0.93	0.93	0.89	33.68	1556.68
TWANG+NN	889.73	1.73	126.27	13.75	191.16	3.29	0.88	1.00	0.93	0.86	14.27	3674.06
TWANG+U+NN	943.90	1.76	72.10	5.16	96.55	3.27	0.93	1.00	0.96	0.92	21.55	2432.89
OTSU+F+NN+U*	989.63	9.84	26.37	0.67	79.67	59.33	0.97	0.99	0.98	0.94	18.66	2809.72

Table 4.8: Quantitative performance assessment of a nearest neighbor tracking algorithm (NN) applied on different segmentation results. Two algorithms without uncertainty-based improvements as described in Sec. 3.3 (OTSUWW, TWANG) were compared to enhanced pipelines that explicitly incorporated prior knowledge-based uncertainty treatment (OTSUWW+U, TWANG+U) as described in Sec. 4.2, Sec. 4.3, Sec. 4.4. Furthermore, an OTSU-based segmentation with additional seed points from LoGNSM+F+U (OTSU+F+NN+U) was used with an adapted tracking algorithm (Sec. 4.5.1). Note that the segmentation produced by OTSU+F+NN+U was not usable for other purposes than tracking due to many merged regions (indicated by (*)). The validation measures correspond to true positives (TP), false positives (FP), false negatives (FN), redundant edges (Red.), missing edges (Miss.) and merged objects (Merg.). Furthermore, recall, precision, F-Score and the TRA measure were calculated as described in App. C. Processing times are average values for applying segmentation and tracking on a single image and were measured in seconds (lower values are better) and voxels per second (higher values are better). For better comparison, the constant processing time offset required for image fusion was skipped here (adapted from [219]).

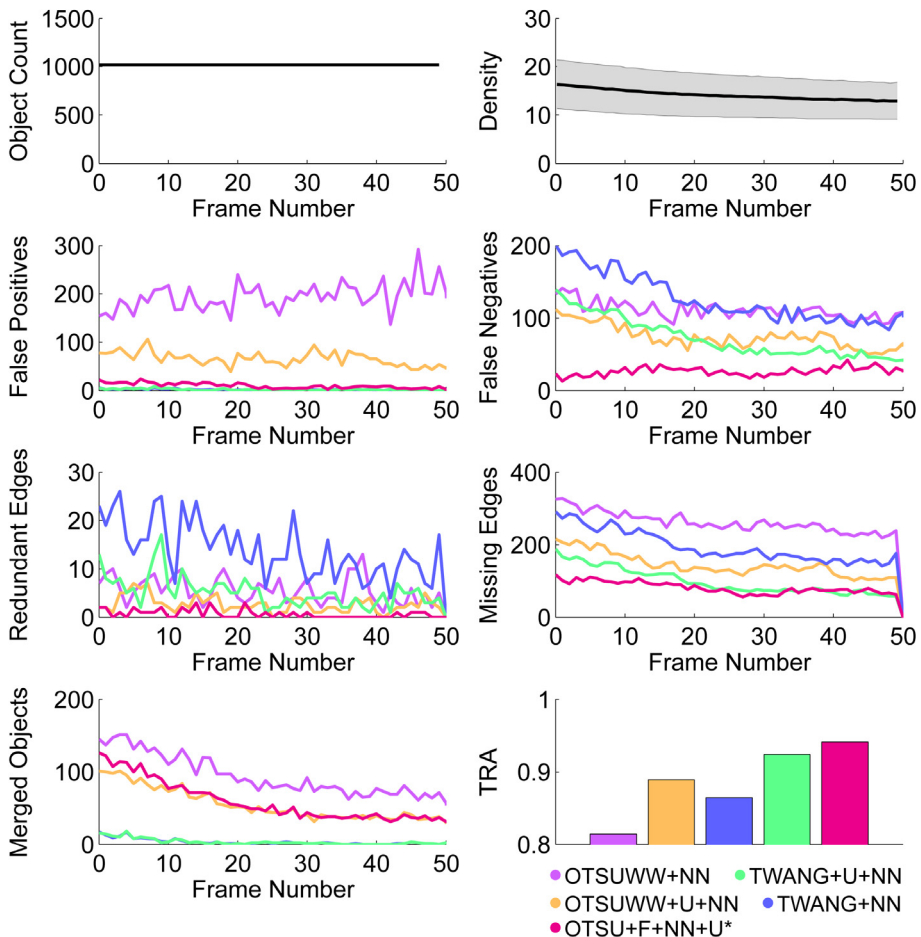


Figure 4.11: Quantitative performance assessment of a nearest neighbor tracking algorithm (NN) applied on different segmentation results obtained on the SBDE4 dataset. Two algorithms without uncertainty-based improvements as described in Sec. 3.3 (OTSUWW, TWANG) were compared to enhanced pipelines that explicitly incorporate prior knowledge-based uncertainty treatment (OUGW+U, TWANG+U) as described in Sec. 4.3. As two rotation images existed per time point (SiMV), the obtained results were fused using the method described in Sec. 4.4. (*) indicates that the respective algorithm did not produce a usable segmentation image, as the correction was solely performed at the tracking step (adapted from [219]).

Thus, an increased amount of false positive detections in background regions was observed for OTSUWW+NN. The reason for the slightly higher number of redundant edges observed for the two TWANG-based methods is not yet fully

clear. Most likely, the seed detection already provided a redundant seed to the segmentation method, which produced two nearby segments that were in turn counted as a redundant segment by the tracking evaluation. However, even for the worst algorithm in this category (TWANG+NN) redundant edges were only observed for 1.4% of the tracked objects and thus play a minor role compared to the other tracking errors. As shown in Fig. 4.11, the number of false negatives and merged objects directly correlated with the density of the objects, *i.e.*, the closer the objects are to each other, the more under-segmentation errors occurred. However, this was not the case for TWANG-based methods due to the explicit prior knowledge about the object size that is incorporated to the algorithm. The OTSU+F+NN+U method produced the best tracking result and was the second fastest method (TRA value of 0.94 and average processing time of 18.7s) closely followed by TWANG+U+NN (TRA value of 0.92 and average processing time of 21.6s). Compared to OTSU+U+NN, TWANG+U+NN and OTSU+F+NN+U provide superior quality in all categories and in particular an increase of the TRA measure by 13.6% and 16.0%, and a decrease of processing times by 58.0% and 63.7%, respectively. Thus, the two latter methods represent the best quality vs. speed trade-off and are suitable for large-scale analyses. Although OTSU+F+NN+U provided excellent results in this comparison, it should be noted that the extracted segmentation masks were largely merged and an object splitting approach as performed for OTSU+U+NN would be required if object properties need to be known. An additional object splitting approach, however, would eradicate the performance benefit of the method and, *e.g.*, the TWANG+U+NN pipeline should be favored in this case.

To fall back on seed point information has no benefit for segmentation methods like TWANG, where the algorithmic design already only extracts a single segment per seed point and literally no merged objects exist. However, an interesting extension to consider in upcoming work might be a combination of the LoGNSM+F+U and the OTSU-based segmentation for seed detection and to feed these seeds to the TWANG algorithm, to reach both a further reduced amount of missed objects and a reduced amount of merged objects. Moreover, the temporal coherence was not yet considered in the investigated framework, *i.e.*, additionally allowing a nearest neighbor matching over multiple frames could potentially also help to reduce the number of missing and redundant detections.

4.6 Discussion

In this chapter, an image analysis pipeline comprised of seed point detection, segmentation, multiview fusion and tracking was systematically equipped with prior knowledge-based enhancements according to the uncertainty framework proposed in Chap. 2. The performance of all improvements was quantitatively assessed on a new and comprehensive validation benchmark inspired by light-sheet microscopy recordings of live specimen. In particular, this framework was successfully employed for the following tasks:

- **Filtering:** The seed detection as well as the segmentation algorithms benefitted from the uncertainty-based feature rejection. Objects that were classified as clearly deviating from the expected object properties were removed and thus the number of false positive detections could effectively be reduced.
- **Splitting:** The segmentation methods based on adaptive thresholding produced many merged segments. Using morphology-based fuzzy sets to describe the properties of a correct detection, objects that were significantly larger could be reliably identified. The improved seed detection methods combined with a watershed-based object splitting successfully separated merged objects.
- **Fusion:** The FSM D values of all objects that were calculated at the seed point detection and the segmentation stage could be used to fuse the information present in complementary views to a single consistent representation of high quality. The feature-based fusion provided almost identical results compared to an intensity-based fusion in the image domain.
- **Tracking:** Similar to the splitting approach, the propagated information was used to resolve tracking conflicts that originated by merged segmentations. By skipping the costly object splitting in the image domain, this approach yielded both improved tracking quality and improved time performance.

For illustration purposes, mostly straightforward processing methods have been used to illustrate the general concepts. However, extending more complex segmentation or tracking algorithms with the presented concepts of filtering, splitting and fusion should work analogously if the uncertainty-based corrections are considered as a post-processing strategy of each processing operator. Especially the tracking step offers a lot of potential to be further improved. The uncertainty framework could be exploited to classify movement events, *e.g.*, to detect object divisions or to reconstruct missing objects using the temporal coherence of the objects. However, this was beyond the scope of this thesis and will be addressed in upcoming work. In addition to the actual algorithmic improvements, it has been shown how the respective fuzzy sets can be parameterized based on available prior knowledge, such as feature histograms or knowledge about acquisition deficiencies. The respective shape of the fuzzy sets has to be determined based on the desired outcome. For instance, the false positive suppression at the seed point detection stage could be performed with a fixed threshold instead of an explicit usage of fuzzy sets for the intensity based features. However, to model the increasing uncertainty (decreasing FSMD value) in regions farther away from the detection objective using a trapezoidal shape was more appropriate. Further work has to be put on the automatic determination of the involved fuzzy sets, *e.g.*, using a semi-automatic approach for a manual classification of a representative subset of data. The remaining sections present the implementations of the methods described so far as well as the biological experiments the methods were applied to.

5 New Implementations and Numerical Tools

After having introduced the theoretical details of the image analysis pipelines considered throughout this thesis, the custom-made software tools that were developed for the practical application of the new methods primarily to large-scale fluorescence microscopy image data are covered in this chapter. The implementations can be separated into XPIWIT (XML Pipeline Wizard for the Insight Toolkit), a C++-based software tool for efficient large-scale image analysis [13] and extensions to the MATLAB toolbox Gait-CAD, which contains numerous tools for sophisticated data analysis [157, 215].

5.1 XPIWIT - XML Pipeline Wizard for the Insight Toolkit

The Insight Toolkit (ITK) offers plenty of features for multidimensional image analysis and has an active community in the biomedical field that constantly improves and extends the functionality of ITK [87]. Current implementations, however, suffer either from flexibility due to specifically compiled C++ pipelines for a certain task or by slow execution times caused by multiple read/write operations for separate filter execution. To overcome these limitations an XML-based wrapper application for the Insight Toolkit has been developed that enables a graphical setup and rapid prototyping of image analysis pipelines while preserving the performance of a pure C++ implementation. Created XML pipelines can directly be used to interface XPIWIT in console mode on large clusters. The current version of XPIWIT already incorporates 79 wrappers to existing ITK filters as well as 27 new filters that were developed in the scope of this thesis. The portfolio of filters comprises IO filters, preprocessing filters, edge detectors, morphological operators, segmentation

filters and many more. To ensure XPIWIT can constantly be extended with new functionality, template files that facilitate the implementation of new filters for XPIWIT have been created. Filters that are already implemented as image-to-image filters in ITK can be included into XPIWIT in a matter of minutes and are instantly accessible via the XML interface. The software tool was successfully applied for the automated analysis of terabyte-scale, time-resolved 3D image data of developing zebrafish embryos.

5.1.1 XML Pipeline Creation

As described in the previous chapters, a general concept of automated image processing, which is also inherent to ITK, is the arrangement of processing operators in a feed-forward pipeline structure. Hence, an XML-based pipeline format has been specifically developed, which allows to create flexible image analysis pipelines using a variety of different processing filters. In Lst. 5.1, a simple example image processing pipeline consisting of an image reader and a median filter is depicted. Besides basic tags about the XML version, encoding and the XPIWIT specifier (Line 1-2), the `<pipeline>` tag encapsulates the actual pipeline. A pipeline in turn is split into separate processing items with the tag `<item>` that can be further customized. Each of the specified filters has a set of inputs that have to point to the item ID of the respective output of a preceding filter. Outputs of the filter are specified by the filter definition in XPIWIT and do not need to be explicitly defined by the user. In addition to image inputs and outputs, XPIWIT also has an internal metadata handling system that can be used to share meta-information such as extracted object locations or image statistics between different processing operators. The optimal execution order of the individual processing operators is internally determined by XPIWIT and ensures that the requested image and meta inputs are available as soon as a processing operator is executed. Moreover, the XML structure allows to adjust all available parameters of each filter using key-value pairs. The XML pipelines are decoupled from the IO parameters, which facilitates the sharing of customized pipelines with other users and additionally allows to use a single pipeline to process multiple images with the same protocol.

```
1 <?xml version="1.0" encoding="UTF-8"?>
  <xpiwit>
3  <pipeline>
    <item item_id="item_0001">
5      <name>ImageReader</name>
      <input>
7          <image item_id_ref="cmd" number_of_output="0">
</input>
9      <arguments>
          <parameter key="SpacingX" value="1.0">
11             <parameter key="SpacingY" value="1.0">
13                 <parameter key="SpacingZ" value="5.0">
15                     <parameter key="MaxThreads" value="12">
17                         <parameter key="BufferOriginalImage"
19                             value="0">
21                                 <parameter key="WriteResult" value="0">
23                                     </parameter>
25                                     </parameter>
27                                     </parameter>
29                                     </parameter>
31                                     </parameter>
25                             </parameter>
27                             </parameter>
29                             </parameter>
31                             </parameter>
          </arguments>
      </item>
    <item item_id="item_0002">
      <name>MedianImageFilter</name>
      <input>
21          <image item_id_ref="item_0001"
23              number_of_output="1">
</input>
25      <arguments>
          <parameter key="WriteResult" value="0">
27             <parameter key="Radius" value="2">
29                 <parameter key="FilterMask3D" value="0">
31                     <parameter key="MaxThreads" value="12">
          </arguments>
      </item>
    </pipeline>
33 </xpiwit>
```

Listing 5.1: Exemplary XML Pipeline.

5.1.2 Data Generation

To apply a predefined XML processing pipeline to a desired set of images, XPIWIT has to be executed from the command prompt with command line input arguments. Alternatively, a configuration text file can be piped to the executable. This configuration file needs to contain the output path, one or more input paths, the path to the XML file to be processed, and further optional parameters. This concept allows to heavily parallelize the data processing, *e.g.*, using distributed computing environments such as Apache Hadoop, where one XPIWIT configuration file is generated per image and all jobs are distributed among available processing nodes (App. B.2). All filters can optionally write their processed output to disk. However, this is mainly used for pipeline debugging and in most cases only the last filter needs to save its output to disk. Metadata that is generated by one of XPIWIT's processing operators is saved in the CSV format (Comma Separated Value). This simple text format can be used to store large tables of values, where each column is separated by a specific character (*e.g.*, "," or ";") and each row is separated by a simple line break. Many powerful data analysis tools (*e.g.*, Excel, MATLAB, Gait-CAD, SPSS, K-nime, RapidMiner) provide functionality to import such CSV files and therefore allow to explore the generated data in more detail. Optionally, XPIWIT is able to store the column specifiers in the first row, which can directly be used, for instance, as identifiers for single features and time series in Gait-CAD.

5.1.3 Special Filters

Besides the various wrappers for ITK-internal filters, the XPIWIT implementation also comprises the following special filters and pipelines:

- Seed detection pipelines (Sec. 3.1)
- TWANG pipeline for efficient segmentation of spherical objects in 3D images (Sec. 3.3)
- RACE pipeline to efficiently segment locally plane-like objects in 3D images (Sec. 3.4)
- Interface to various noise and point spread function simulations for artificial image disruption (Sec. 3.1, Sec. 3.3, Sec. 3.4 and Sec. 4.1)

- Utility filters for converting label images to feature maps (Fig. 4.9)
- Filter for watershed-based object splitting (Sec. 4.3.2)

5.1.4 Gait-CAD Compatibility

The main intention of XPIWIT was to speed up time costly processing steps that directly operate on large-scale image data. However, to exploit the powerful data analysis capabilities of the open-source MATLAB toolbox Gait-CAD [157, 215], it is desirable to have a convenient interface between Gait-CAD and XPIWIT. There are in principal two ways to accomplish the synergy of both software tools. The first option is to extract data from large-scale image data solely using XPIWIT and to perform the analysis of extracted information by importing the generated `*.csv` metadata files into Gait-CAD data structures. As a second option, all filters that are compiled into the respective XPIWIT executable can be used within the usual Gait-CAD plugin-sequence dialog. Gait-CAD automatically takes care of generating temporary input files and XML pipelines for XPIWIT and automatically imports the results generated by XPIWIT back to Gait-CAD. However, the general concept of pipelines in Gait-CAD assumes strictly linear pipelines with a single predecessor and a single successor. To achieve more flexibility and to be able to exploit the entire functionality of XPIWIT, a dedicated GUI has been developed, which is described in more detail in the following section [84].

5.1.5 Graphical User Interface for Rapid Prototyping

An easy-to-use graphical user interface (GUI) has been developed that allows to create and modify XML pipelines based on the filters compiled into the XPIWIT executable. In Fig. 5.1, the basic layout of the GUI is depicted. Filters can be placed via drag and drop from the filter list (1) to the pipeline drawing area (2). All parameters can be customized by selecting the respective filter in the drawing area and by adjusting the parameters in the parameter customization panel (3). The GUI is implemented using the open-source version of the Qt framework, which allows us to compile the GUI for Windows, Linux and Max

OS X. Once the input and output parameters (4) are defined, the created pipeline can immediately be run on the specified data and can of course be saved and re-opened using the menu bar (5).

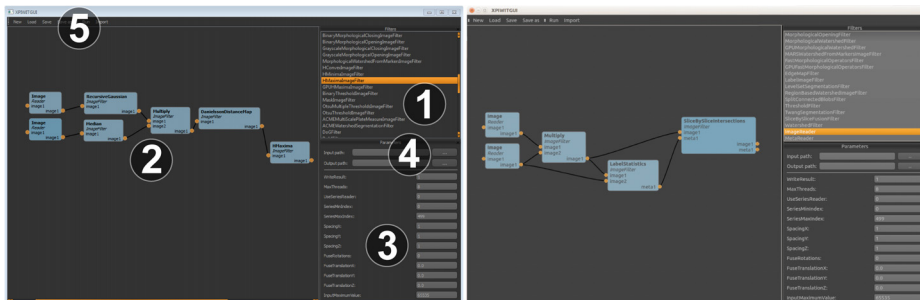


Figure 5.1: Exemplary screenshot of the XPIWIT GUI on Windows (left) and on Ubuntu (right). The main control elements of the GUI are the filter list (1), the pipeline drawing area (2), the parameter adjustment panel (3), the input/output definition panel used for pipeline test runs (4) and the usual controls for saving, loading and executing pipelines (5). Generated pipelines are independent from the GUI, *e.g.*, to perform large-scale analyses on computing clusters.

5.1.6 Implementation Details

XPIWIT is implemented in C++ using the Insight Toolkit (<http://www.itk.org/>) and the Qt SDK (<http://qt-project.org/downloads>). Platform independent project files have been realized using the CMake build tool (<http://www.cmake.org/>). Both, XPIWIT and the associated GUI have been successfully compiled and tested under Windows, Linux and Mac OS X. Software and documentation are licensed under the GNU-GPL and will be made publicly available for download under <https://bitbucket.org/jstegmaier/xpiwit> as soon as the corresponding manuscript is published.

5.1.7 Comparison to Existing Software Solutions

To compare XPIWIT to other existing software tools, Tab. 5.1 shows an extended version of the software comparison presented by Kankaanpää *et al.* [92]. Three new categories that are crucial for the analysis of large-scale imaging experiments were added. Namely, algorithms for large-scale analyses need to be

computationally efficient to guarantee reasonable processing times (*Fast*) and they have to be implemented properly to make the applications usable with possible memory constraints (*Memory Efficiency*). Additionally, the usage of distributed processing on computing clusters often requires the applications to be runnable in headless mode and parametrizable via command line (*Parallelizable*). Besides the new evaluation criteria, two new applications (ICY [41] and XPIWIT [13]) were included in the table. Commercial software tools (Imaris, FluoView, Volocity and ZEN Lite) and applications that are solely suitable for visualization or other special purposes (BioImage Suite, VisBio) were excluded from the overview table. The performance estimation with respect to speed and memory efficiency was performed by applying both a median and a Gaussian filter to a large image dataset.

Name	Open	Ext.	Usable	Adj.	Appl.	Exp.	Fast	Mem.	Parall.
BioImageXD	••••	••••	••••	••••	••••	••••	••	•	•
Fiji	••••	••	••	•••	••	••••	••	••••	••••
Vaa3D	•••	••	•••	•••	•	••••	••••	••••	••••
Matlab	••	••	•••	••••	•	••	••	•	•••
ICY	••••	••••	•••	••••	•••	••••	•	•	•
<i>XPIWIT</i>	••••	••	•••	•••	••••	••••	••••	••••	••••

Table 5.1: Comparison of the software tools BioImageXD [92], Fiji [200], Vaa3D [174], Matlab (The MathWorks, Inc.), ICY [41] and XPIWIT [13]. As described in [92], the evaluation criteria are open-source (*Open*), extensibility (*Ext.*), usability (*Usable*), adjustability (*Adj.*), applicability (*Appl.*), expandability (*Exp.*), processing speed (*Fast*), memory efficiency (*Mem.*) and the possibility to run the software on distributed computing environments (*Parall.*). The properties are rated from one to maximally four points. For the criteria *Fast* and *Mem.* processing times and memory consumption were measured for a median and a Gaussian filter applied on a large image dataset. It should be noted that all assessments are of objective nature and each of the applications has its own benefits and drawbacks in specific fields of application (extended from [92], additional table entries are highlighted in italic font).

As indicated in Tab. 5.1, the strengths of XPIWIT clearly lie in the domain of efficiently processing large-scale data, potentially on computing clusters. Of course, these criteria can also be achieved with special purpose implementations such as [5], but XPIWIT additionally provides the flexibility of using different pipelines with a single static executable. Further work for XPIWIT has to be put into adding more wrappers to ITK internal filters as well as further improvements of the GUI with respect to result previewing, parameter adjust-

ments and the possibility to use VTK-based visualizations. With respect to the new measurements, Vaa3D was the most powerful alternative to XPIWIT. As both XPIWIT and Vaa3D build on the foundation of ITK, the processing times were mostly comparable. However, Vaa3D encapsulates a lot of functionality into black box operators that lack parameter adjustment possibilities, *i.e.*, the algorithmic behavior is not always transparent to the user.

5.2 Extensions of the Open-Source MATLAB Toolbox Gait-CAD

Gait-CAD is a convenient open-source MATLAB toolbox that is being developed at the Institute for Applied Computer Science at KIT since 2006. It is especially suited for interactive data analysis of both single features and time series [157]. To extend the functionality of Gait-CAD with respect to the automated analysis data extracted from large-scale images, several additional toolboxes have been developed to facilitate image analysis, tracking and benchmarking as well as the visualization and export of related experimental results. The following sections provide a brief overview of the individual toolboxes and a more detailed description can be found in [215, 221]. The most recent releases of Gait-CAD including the ImVID and the Tracking extension packages can be obtained from <http://sourceforge.net/projects/gait-cad/>. Furthermore, the SpinalCord extension package is hosted at <http://sourceforge.net/projects/zebrafishimage/files/>. The remaining extension packages are still under active development and the current versions are available upon demand.

5.2.1 The ImVID Extension Package

The ImVID Extension Package extends the powerful data mining capabilities of Gait-CAD by comprehensive functionality for automatic image processing and image analysis. To facilitate the input and output of the toolbox and to make it accessible to a broad range of scientific images the BioFormats library is used [129]. Imported image data is immediately linked with Gait-CAD's internal data structures, such that analyses can always benefit from both data

and visualizations as well as raw and processed images. The toolbox provides convenient access to use and parameterize numerous filters of the MATLAB Image Processing Toolbox as well as custom-made filters such as the methods described in Sec. 3.1 and Sec. 3.3 as well as all Gait-CAD compatible XPIWIT filters that map a single input image to a single output image. All implemented filters can be used and combined to form complex image analysis pipelines using Gait-CAD's plugin sequence functionality. Once a pipeline is created it can be easily applied to all imported images automatically. The extension package can also be used in combination with Gait-CAD's batch processing system to easily scale projects up, *e.g.*, to analyze thousands of images of a high-throughput screen in an automatic fashion.

5.2.2 The Tracking Extension Package

The Tracking Extension Package can be used to identify corresponding objects in multidimensional time series of object locations. Based on available prior knowledge about the objects' movement behavior, several constraints such as bounding box size and maximum velocity can be specified. Objects missing in a single frame that reappear in the next frame can be reconstructed by a linear interpolation scheme. Furthermore, a heuristic was implemented to repair incomplete tracks that originate from objects missing in multiple frames or by events such as object divisions. Tracked data can either be viewed using Gait-CAD's internal visualization functionality or exported in the `*.vtk` format to generate sophisticated videos of the object tracking, *e.g.*, in Kitware's scientific visualization tool ParaView [80]. The toolbox was successfully applied to the automated analysis of embryonic development as described in [106, 227, 244] and in Sec. 6.2.

5.2.3 The SpinalCord Extension Package

The SpinalCord Extension Package basically condenses all functionality to perform a quantitative analysis of laterally oriented zebrafish embryos in fluorescence microscopy images using the methodology described in Sec. 3.2. It is comprised of components for

- splitting and merging RGB channels

- rotation correction of the images
- automatic focus detection and extended focus image generation
- detection, extraction and transformation of elongated objects
- extraction of intensity profiles
- extraction of object distribution histograms.

Although the implementation was specifically designed for the particular purpose of automatically analyzing the spinal cord region of laterally oriented zebrafish embryos, the modular design in combination with all available ImVID filters allows a straightforward adaption to similar problems, where the main task is to detect elongated objects (possibly with colocalized information in multiple color channels) and to transform the extracted regions to a uniform and comparable representation. The method was successfully applied to the automated analysis of the impact of small molecules and the effect of environmental Acetylcholinesterase (AChE) inhibitors to particular neuron populations in specifically designed transgenic zebrafish embryos [204, 221].

5.2.4 The Embryo3DT Extension Package

The Embryo3DT Extension Package comprises numerous functions and tools that facilitate the automated analysis of data extracted from large-scale microscopy images using XPIWIT. The following functionality has been implemented:

- a registration algorithm for the fusion of extracted spatial information of dynamic objects in different rotation images (Sec. 4.4)
- functions to reject false positive detections based on fuzzy set membership values (Sec. 4.2)
- fusion heuristics to get rid of redundant detections (Sec. 4.2.1)
- multiple specialized visualization options to optimally view the extracted data.

In Fig. 5.2, an exemplary visualization for the interactive analysis of embryonic development is shown. Besides the possibility to instantly visualize the dynamics of extracted objects, the Embryo3DT Extension Toolbox can export object locations and associated feature values to a `*.vtk`-based format to exploit the interactive visualization capabilities of ParaView.

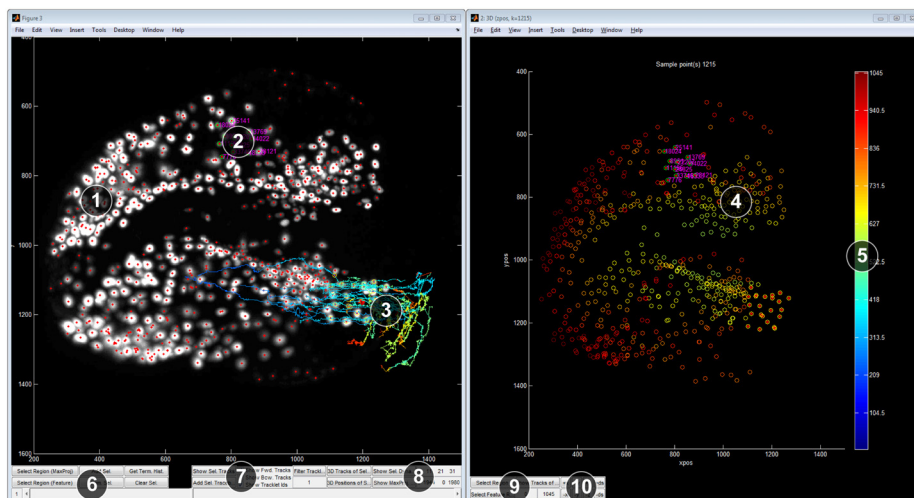


Figure 5.2: Exemplary screenshot of an Embryo3DT visualization that shows extracted objects superimposed on the respective maximum intensity projection (1) with associated tracklet IDs (2) and a group of tracked objects (3). The right panel depicts a color-coded feature visualization (4, 5) that can be used for selection and tracking visualization in the feature space. The control elements are used to select data and to navigate in the temporal domain (6), to visualize tracks of selected objects (7), to generate object videos (8), to perform feature based selections and visualizations (9) and to manually refine the registration of complementary rotation views (10).

5.2.5 The Benchmark Extension Package

The Benchmark Extension Toolbox includes all the functionality for generating a simulated benchmark of moving and dividing objects as described in Sec. 4.1. All parameters can be easily adjusted using a dedicated Gait-CAD dialog. It can be used to generate simulated 2D and 3D images of time series of dynamic objects or independent frames with a fixed number of objects. Next to the simulated image, segmentation label images as well as object metadata are stored

as ground truth, to enable subsequent evaluation and comparison of automatically obtained processing quality of different image analysis operators. The toolbox comes with a database of simulated fluorescent nuclei that were extracted from the data generated by Svoboda *et al.* [226]. The artificial images can be distorted using different noise models, such as Poisson noise, Gaussian noise or salt and pepper noise and can be convolved with an arbitrary point spread function [75]. As MATLAB still tends to be slow when working with large 3D images, the acquisition simulation was additionally implemented in XPIWIT, to speed up the process for the creation of large-scale benchmarks as discussed in Sec. 4.1. In addition to the generation of simulated objects, functions to quantitatively compare the ground truth information to the respective results of automatic analysis operators are provided (Sec. 4.1.3).

5.2.6 Semi-Automatic Uncertainty-based Image Analysis

Although a high degree of automation is indispensable for large-scale image analysis projects, in some cases it makes sense to proofread and curate extracted information manually. For this purpose, a semi-automatic procedure has been developed that starts with a very sensitive automatic detection using XPIWIT and proceeds with a manual post-processing step to refine the results produced by the automatic processing. This allows to process large datasets in one run with potentially sub-optimal parameters and to correct the results afterwards instead of performing all calculations over and over again until optimal parameters are found. The developed GUI is perfectly suited to adjust spatial constraints and intensity thresholds to largely suppress false positive detections. Moreover, it can be used to estimate prior knowledge-based fuzzy set membership functions for all available features. Using an intuitive mouse wheel-based adjustment of a selected feature threshold, an instant visual feedback of filtered results is provided to the user and the respective threshold values can be used to set the parameters of a trapezoidal membership function (see (4) in Fig. 5.3) as well as to identify the uncertainty framework filter parameters α_{il} and β_{il} (Sec. 2.4). The individual objects are color-coded in red, blue and green to visualize low, medium and high degree of membership to the current fuzzy set. The GUI allows navigating through all files of a project, allows to propagate parameter settings to all other time points or to adjust all

parameters individually. Filtered data can finally be exported to a Gait-CAD compatible `*.csv` file. To enable browsing through large datasets, internally only file references are stored and images are loaded on the fly. Fig. 5.3 shows an exemplary screenshot of the GUI where a 3D benchmark image as discussed in Sec. 4.1 is analyzed. The semi-automatic tool was used for counting motoneurons in zebrafish embryos at different developmental stages and to refine segmentation results obtained from large-scale analyses of embryonic development that were extracted using XPIWIT.

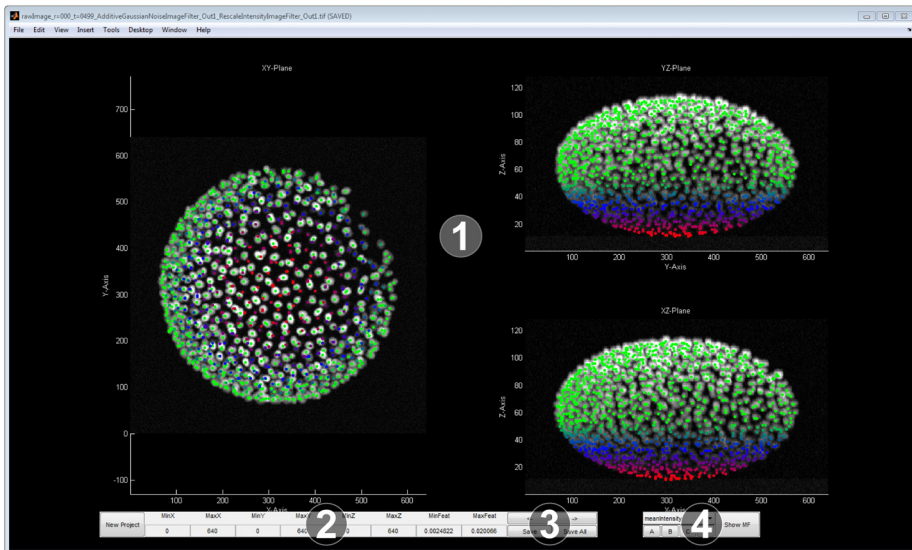


Figure 5.3: Exemplary screenshot of the semi-automatic data correction GUI. The three main panels depict the maximum intensity projection views of the XY, XZ and YZ plane of an image containing simulated cell nuclei (white objects) with superimposed centroids of detected objects (1). Undesired objects can be filtered based on the feature values, *i.e.*, spatial constraints or intensity constraints (2). Parameters can be propagated to all other images of the project and filtered results can be saved as Gait-CAD compatible `*.csv` files (3). Parameters used in the uncertainty framework, *i.e.*, the fuzzy set membership parameterization of any of the available features can be performed in this GUI. The color-code from red (low FSMD) over blue (medium FSMD) to green (high FSMD) provides an instant visual feedback of the membership degree obtained for the individual objects (4).

6 Automated Quantitative Analysis of Embryonic Development in Large-scale Fluorescence Microscopy Images

After the new methodology and the associated implementations have been presented in the previous parts of the thesis, this chapter introduces three application examples of the developed methods that all stem from the field of developmental biology. New imaging techniques such as light-sheet microscopy [85, 97, 230] as well as new methods to automate image-based experiments [176] offer unprecedented possibilities for detailed new studies in this scientific field. However, the immensely large amount of image data that is generated routinely by these novel techniques has to be stored, processed and analyzed with a high degree of automation as well to keep pace with the fast acquisition and to overcome the limitations of manual investigations. The challenges observed in the presented examples are a large amount of multi-channel 2D images that needed to be quantitatively compared to each other (Sec. 6.1) and terabyte-scale time series of 3D images produced by state-of-the-art light-sheet microscopes (Sec. 6.2, Sec. 6.3).

6.1 Automated Quantification of Neuronal Patterns in the Spinal Cord of Zebrafish

The aim of this project was to automatically quantify the impact of unknown chemicals to different types of neurons in the spinal cord of laterally oriented zebrafish embryos [221]. The analysis was performed on a newly established transgenic line of zebrafish that exhibited fluorescence upon gene expression of *dbx1b* (dI6, V0 and V1 interneurons tagged with eGFP) and *ngn1* (Rohon-Beard

cells, motoneurons, tagged with RFP) at two distinct excitation wavelengths [221]. The advantage of using this particular double-transgenic zebrafish line in the experiment was to have as many as possible types of central nervous system neurons (sensory-, inter- and motoneurons) for quantitative assessment in a single biological system, to possibly identify and thus narrow down the interesting candidates which may have effects on spinal cord neurons upon treatment. Treatment effects could for example be an increase or decrease in the number of a specific neuron type or an observable positional shift of the investigated cells. To enable high-throughput screening experiments where thousands of embryos might be screened, the developed algorithms should provide a strongly automated readout and extract both quantitative and qualitative information reliably and fast. As there was no existing solution to the specific problems observed for this screen, a new image analysis pipeline had to be developed (Sec. 3.2). Here, a pilot study to confirm the applicability of the proposed analysis framework is presented. Therefore, a systematic validation was performed using embryos that were treated with different concentrations (0.5 – 10 μM) of the Notch inhibitor LY-411575 with a known effect on the development of moto- and interneurons [60, 206].

6.1.1 Dataset Description

Images were acquired using compound fluorescence microscopy (Leica DM5000 at 10 \times magnification). The embryos were excited with two different wavelengths (GFP: 484 nm, RFP: 585 nm) and the emitted fluorescence was captured using RFP, GFP and CY3 filters. The CY3 filter was used to facilitate the acquisition procedure, as the RFP signal was also visible in this channel due to fluorescent crosstalk and the additional channel allowed to use a different exposure times for the CY3 channel to make the dimmer motoneuron signal sufficiently visible (RFP: 40 ms, GFP: 200 ms, CY3: 410 ms). The considered images were acquired using a dynamic range of 8 bit at a resolution of 1392 \times 1040 pixels (px) with a physical pixel size of 1px $\hat{=}$ 0.64 μm . An overview of the acquired raw image data is shown in Fig. 6.1A. In-focus images were selected manually to have three high-quality images per concentration and repeat, which yielded images of 108 embryos (3 repeats, 12 different concentrations and 3 embryos per concentration). Further details on the sample preparation, microscopy parameters and the experimental procedure can be found in [221].

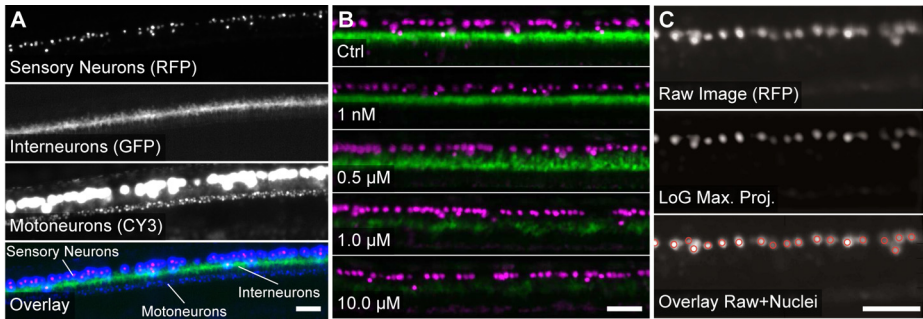


Figure 6.1: (A) Raw images of the sensory neurons (RFP channel), the interneurons (GFP channel), the motoneurons (CY3 channel) and an overlay of all three channels. (B) Exemplary straightened crop regions obtained by the ROI extraction algorithm described in Sec. 3.2.2. The overlay images of the RFP and the GFP channel are shown for a control embryo and embryos treated with different concentrations of LY-411575. (C) Exemplary results of the LoG-based seed detection applied on the straightened ROI images (algorithm from Sec. 3.1.2). Scale bar: $50 \mu\text{m}$ (adapted from [221]).

6.1.2 Automated Image Analysis Framework

The automated analysis framework was developed using as much prior knowledge as possible. In particular, the interneuron signal (GFP channel) was used to estimate the location of the spinal cord. The other neuron populations were localized relative to this identified reference line, *i.e.*, sensory neurons and motoneurons were detected above and below the spinal cord center line, respectively. A more detailed explanation of the specific properties of the individual acquired channels is summarized Tab. 6.1. The first step of the analysis procedure was the extraction of a comparable region of interest of the embryos. The interneuron signal appeared as a dense elongated object along the spinal cord of the embryo and was extracted using the method described in Sec. 3.2. Based on the determined spinal cord in the GFP channel, the same crop region along the regression curve with radius $r = 70$ was extracted from all channels, to obtain a colocalized straightened version of the region of interest (Sec. 3.2.2). To identify sensory neurons and motoneurons in the RFP and the CY3 channel, respectively, the LoG-based seed detection method was used with $\sigma_{\min} = \{2, 3\}$, $\sigma_{\max} = \{3, 4\}$ and $\sigma_{\text{step}} = \{1, 1\}$, where the parameters were determined on the basis of minimum and maximum radii of the respec-

tive neurons (Tab. 6.1) as described in Sec. 3.1.2. For each of the channels, all single features described in Sec. 3.2 were extracted and used for further analysis and comparison. Exemplary overlay images of straightened crop regions for different concentrations as well as detected sensory neurons in an RFP channel image are shown in Fig. 6.1B, C. The performance of the region of interest extraction and the LoG-based neuron detection was validated as described in [221]. The software required less than 6 seconds per image on a single processor of a common desktop computer, *i.e.*, even larger screens comprised of several thousands of images can be processed within a few hours.

Type	Channel	Description
Inter-neurons	GFP	Thick band of cells between sensory- and motoneurons, individual cells not detectable, cubic polynomial to model the spinal cord
Sensory Neurons	RFP, CY3	Mostly well separated nuclei on dorsal side of the spinal cord, mainly overexposed in CY3, good signal in RFP channel, diameter: 7 – 10 μm
Moto-neurons	RFP, CY3	Densely packed nuclei on the ventral side of the spinal cord, weak signal in RFP channel, good signal in CY3 channel, diameter: 5 – 7 μm

Table 6.1: Prior knowledge used for the algorithmic design. Individual channels were perfectly colocalized and exhibited the fluorescent signal of stained nuclei from three different neuron types with channel-specific exposure times and parameters as described in Sec. 6.1.1 (adapted from [221]).

6.1.3 Results

On the basis of visual inspection and the results reported in [60, 206], a constant number of sensory neurons, a decreased GFP intensity of the interneurons and an increase in the number and intensity of motoneurons upon treatment with various concentrations of LY-411575 was anticipated. This was confirmed by the performed dose-response experiment and is exemplarily summarized in Fig. 6.2 for the GFP signal of the interneuron population.




Figure 6.2: (A) The dose-response curve of the GFP peak intensity (pvp_{GFP}) shows an exponential decrease for increased concentration of LY-411575. (B) GFP channel intensity profiles of the pilot screen classified into *no effect* (red), *slight effect* (green) and *strong effect* (blue). The interneuron intensity was slightly reduced for a concentration of 0.5 μM and was significantly reduced by higher concentrations of LY-411575. The gray surface represents the norm corridor of one standard deviation around the average control embryo profile (Sec. 3.2.2). (C) Color-coded intensity profiles of the GFP channel. Each row of pixels corresponds to a single intensity profile and the curves are sorted by increasing LY-411575 concentration (bottom to top). In (B) and (C), a shift of the interneuron peak location towards the ventral side of the embryo can be observed for a treatment of 0.5 μM . This effect can also be visually observed Fig. 6.1B. (adapted from [221]).

As suspected, the GFP intensity of interneurons was significantly decreased upon LY-411575 treatment compared to the control embryos. The effect was reflected in all GFP related features (e.g. i_{GFP} , pvp_{GFP} or nd_{GFP}). The dose-response curve depicted in Fig. 6.2A shows an exponentially decreasing GFP peak intensity value for increased treatment concentrations. Additionally, a tendency to an increased number of motoneurons could be observed in the CY3 channel for LY-411575 concentrations larger than 4.0 μM , e.g., using the features pvp_{CY3} or nd_{CY3} as described in [221]. One possibility to visualize the extracted information of the entire experiment in a single plot is the superposition of all intensity profiles as depicted in Fig. 6.2B. The shaded area in Fig. 6.2B indicates the norm corridor of one standard deviation. Profiles that clearly deviate from this area may indicate a potential treatment effect and should be investigated in more detail. In the case of the investigated LY-411575 concentrations, for instance, the reduction of interneuron GFP intensity could roughly be separated into the classes *no effect*, *slight effect* and *strong effect* for control embryos, a concentration of 0.5 μM and larger concentrations, respectively (Fig. 6.2B).

The decreased GFP intensity was caused by a decreased number of interneurons. As the counting of interneurons was not possible using the considered image material, this was confirmed using 3D image stacks acquired with confocal microscopy as described in [221]. Another powerful visualization of the extracted information is the side-by-side alignment of the intensity profiles of the GFP channel images as shown in Fig. 6.2C. Using this visualization a comprehensive comparison of all intensity profiles can be obtained within just one overview image. Effects like a ventrally oriented shift of the interneurons for a concentration of $0.5 \mu\text{M}$ could be easily identified with this approach. Further results obtained on the remaining image channels and the associated features that were used to quantify the effects can be found in [221].

6.1.4 Discussion

Various methods for the automated analysis of zebrafish related screens were presented in the past [44, 132, 133, 158]. However, the specific demands and challenges of the presented experiment could not be handled by any existing methods and therefore needed to be newly developed. The generally introduced method for the extraction of elongated objects and their environment presented in Sec. 3.2 was inspired by the demands of this screen and worked well in practice as confirmed by the results presented in Sec. 6.1.3. Besides the validation performed on the simulated benchmark in Sec. 3.2, manual investigation confirmed that the algorithm properly extracted 98% of the 108 images in the pilot screen using the regression-based ROI extraction method and thus represents a robust solution to this specific problem. Additionally, $90.0 \pm 3.3\%$ of manually annotated sensory neurons were successfully detected by the LoG-based approach confirming its validity and sufficient quality to detect trends in changed sensory neuron counts. It should be noted that even manual counting was ambiguous in some cases and intra- and inter-expert variations were observed. The resolution used in the pilot screen was too low to unambiguously resolve the individual motoneurons. Thus, to quantify the effects on the motoneuron development it was more reasonable to use the intensity based features such as intensity profiles, spatial positioning of intensity peaks with respect to the spinal cord center or the spatial distribution analysis approaches

as described in Sec. 3.2.2 and [221]. Alternatively, a higher magnification level of the microscope or even a different microscopy technique could potentially improve the motoneuron detection as detailed in [221].

Overall, the presented approach represents a well-suited method to quantitatively analyze different types of neurons in the spinal cord of zebrafish. The versatility of the extracted features could be used to reliably characterize treatment effects in a reproducible manner and enabled the extraction of image properties on the pixel or intensity level that were beyond visual inspection capabilities of human experts. Besides the application on the described pilot screen, the proposed image analysis framework was already successfully applied to analyze a small molecules screen of a large chemical library consisting of about 1600 chemicals with unknown toxicological effects. Furthermore, with minor parameter adjustments, the presented method was applied to quantitatively assess the effect of environmental Acetylcholinesterase (AChE) inhibitors in a specifically designed transgenic zebrafish line that served as a fluorescent biosensor [204].

6.2 Automated Quantitative Analysis of Embryonic Development in 3D+T Microscopy Images

As sketched in the introduction, light-sheet fluorescence microscopy offers a great potential for in-depth analyses of living specimens such as the fruit fly, zebrafish or mouse embryos at the single cell level [99, 220]. The ultimate goal would be to reconstruct an accurate lineage tree of cellular ancestry starting at the fertilized egg and going down to the various differentiated cell types of a living organism [5, 227]. The resolution and image quality achieved by light-sheet microscopy allows to detect and quantify phenotypic alterations in mutant embryos or upon drug exposure at the single cell level [106, 170]. However, the large amount of acquired image data cannot be manually analyzed and requires sophisticated and fast algorithms to exhaustively explore the recorded data. Although many methods for the automatic analysis of such datasets have already been presented, most of them were developed for 2D applications and do not scale properly. Thus, the applicability of most existing methods to large-scale datasets routinely acquired in developmental biology imaging experiments is rather limited (Sec. 1.1.2).

In this section, image datasets acquired using state-of-the-art light-sheet microscopy and the developed pipeline for their automatic analysis are introduced. While the new imaging technique allows to capture detailed 3D+t images of the early embryonic development of zebrafish embryos, there were multiple challenges that had to be solved to automatically extract biologically relevant data accurately and fast. The theoretically introduced methods derived in the last chapters were applied on these large-scale datasets to demonstrate the successful operation for seed detection, segmentation, multiview fusion and tracking on an actual biological problem and to emphasize the potential of quantitative approaches in developmental biology.

6.2.1 Dataset Description

The image data used for the automated analysis of early embryonic development in zebrafish as described in this chapter were acquired using a custom-made digital scanned laser light-sheet fluorescence microscope (DSLIM) [106, 170] (Sec. 1.1.1). The microscope was equipped with two-sided illumination

and a rotatable sample chamber that allowed a rotation by 180° around the vertical axis. Dechorionated transgenic zebrafish embryos with fluorescently labeled nuclei were embedded in low melting agarose, which was enriched with fluorescent polystyrene beads with a diameter of $0.1 \mu\text{m}$ that served as landmarks for the rigid registration of the different view angles. In a typical experiment for the analysis of an early zebrafish embryo, the image acquisition was started between 1-2 hours post fertilization (hpf, 8-64 cell stage) and continuous imaging was performed for 10-16 hours in one minute time intervals (about 30 seconds for a single view). A single image stack acquired using this procedure had a resolution of $2560 \times 2160 \times 500$ voxels, a dynamic range of 12 bit (stored as 16 bit unsigned char), a file size of about 5.16 GB and covered a physical volume of roughly 1 mm^3 . Thus, the amount of image data that had to be analyzed was on the order of 6-10 TB for a single embryo experiment. Several exemplary snapshots and maximum intensity projections of different views of a developing zebrafish are shown in Fig. 6.3A and Fig. 6.3B, respectively. Further information on the microscopy setup, the experimental procedure and the animal handling are provided in [106, 170]. Information on the infrastructure for storing and processing the large-scale datasets is summarized in App. B.

6.2.2 Automated Image Analysis Framework

In addition to the terabyte-scale dataset size, the main problems observed during the automated analysis were the reduced axial resolution (Fig. 6.3B) as well as light scattering and intensity reduction in regions farther away from the detection objective (Fig. 6.3B). The employed rotation of the specimen at every second frame helped to obtain a reasonable image quality of the entire embryo (analogous to the SeMV images used in Sec. 4.4). However, the use of this multiview acquisition required an additional fusion step to combine the extracted information of two complementary views into a single high-quality representation. Besides the problem of determining the exact transformation used for the registration of the two views, the dynamics of the investigated specimen had to be considered, *e.g.*, if investigated objects were moving during the acquisition of the complementary views (Fig. 6.3C).




Figure 6.3: (A) Snapshots of early embryonic development in zebrafish (4–10 hpf) using an axial maximum intensity projection (MIP). (B) MIPs of two complementary rotation images projected along the z , y and x direction, respectively (12 hpf). (C) Slices of two complementary views that have been registered using a bead-based registration approach. Displacement of cyan vs. magenta objects indicates cell movement happening between two frames. Scale bar: 50 μm .

Multiview Information Fusion

The identification of the rigid transformation between the complementary rotation views was accomplished using a bead-based registration algorithm similar to the one proposed by Preibisch *et al.* [185]. The fluorescent beads were detected using the LoG-based seed detection method as described in Sec. 3.1. Correspondences between detected beads were identified by a matching of descriptors that characterized the local relations of each bead to its direct neighbors [185]. The iterative RANSAC optimization algorithm used in [185] was replaced by an analytic solution to the least-squares estimation problem based on the SVD-based (Singular Value Decomposition) method proposed by Umeyama *et al.* [232], to speed up the transformation calculation. In Fig. 6.4, the detected beads, the identified correspondences and the final registration result are shown.

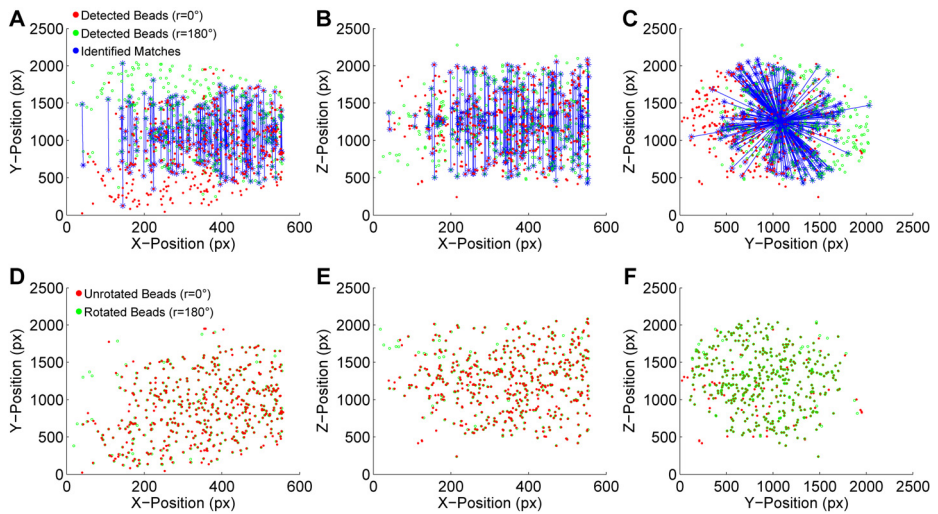


Figure 6.4: Registration of beads detected in the Laplacian-of-Gaussian maximum projection using the parameters $\sigma_{\min} = 4$, $\sigma_{\max} = 5$, $\sigma_{\text{step}} = 1$. As the beads are smaller than the nuclei, the LoG parameters were adjusted to match the manually measured radii of the beads (Sec. 3.1). The first row (A-C), shows the unregistered beads and the identified correspondences of the seeds detected in two complementary rotation images. The second row shows the registered beads (D-F). The average distance of registered beads was $3.07 \pm 1.43\text{px}$ in image coordinates, which corresponds to $1.22 \pm 0.58\mu\text{m}$ in physical coordinates.

If the complementary images were acquired simultaneously or if the investigated specimen was fixed, the identified transformation can be used to directly fuse the views to a single high-quality image. For the image material considered here, however, a direct fusion of the acquired rotation images to a single image stack did not lead to satisfactory results. This was caused primarily by object movements between rotationally complement frames and reduced image quality in the central regions of the fusion images (Fig. 6.3C). Thus, pre-processing, seed detection and segmentation were performed independently for each of the rotation images and then fused to a single representation of the entire embryo using the identified rigid transformation. To eliminate redundant information, *i.e.*, objects that were detected in both of the views, the information fusion scheme described in Sec. 4.4 was used.

Fast Detection and Segmentation of Fluorescently Labeled Nuclei

The seed detection was performed using the LoG-based approach as described in Sec. 3.1 and [220]. The images were down-scaled by a factor of 0.5 in the xy-direction, which still provided sufficiently high resolution to set individual nuclei apart from each other and helped to reduce the required processing time. The LoG parameters ($\sigma_{\min} = 6$, $\sigma_{\max} = 9$ and $\sigma_{\text{step}} = 3$) were determined using manually measured minimum and maximum radii of the imaged nuclei as described in Sec. 3.1. Based on the extracted seed points, the fluorescently labeled nuclei were segmented using the TWANG method with the default parameters $\sigma_{\text{grad}} = 3.0$, $\sigma_{\text{kernel}} = 3.0$ and $\omega_{\text{kpm}} = 1.5$ as described in Sec. 3.3 and [220]. To save both time and storage space, the actual segmentation images were not written to disk and only the extracted properties of detected nuclei were stored in a table format and used for further processing.

The extracted information of complementary views was registered using the identified transformation of the bead-based registration and then fused to a single representation (see Sec. 4.4 and [106]). To eliminate redundant objects, each detected nucleus was associated with a fuzzy set membership degree of being a *desired object/nucleus* (high value) or a *false positive/noise* (low value). The most discriminative features for the extracted segments were identified by manual inspection of the feature histograms at different developmental time points. As the fluorescent signal was not yet strongly expressed at the early developmental stages, the algorithm detected mostly false positives, which could

be used to characterize the expected feature values of false positive detections. Contrary to this, later time points contained almost no false positives and were thus used to determine feature values of true positive detections. The object volume, the minor axis, the ratio of foreground vs. background volume and the foreground vs. background intensity ratio of each nucleus were identified to be the most discriminative features. A visualization of the corresponding fuzzy set membership functions $\mu_{\text{vol}}(x)$, $\mu_{\text{mia}}(x)$, $\mu_{\text{fg/bg, vol}}(x)$, $\mu_{\text{fg/bg, int}}(x)$ as well as the respective parameter vectors are provided in Fig. 6.5A-D.

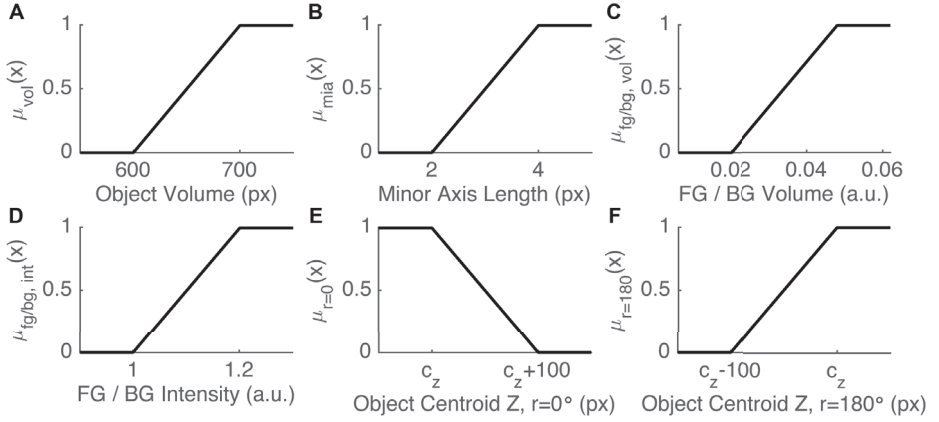


Figure 6.5: Fuzzy set membership functions that were used to discard false positive detections during the multiview fusion. Panels (A-D) show the membership functions for the object volume ($\theta_{\text{vol}} = (600, 700, \infty, \infty)^\top$), the minor axis ($\theta_{\text{mia}} = (2, 4, \infty, \infty)^\top$), the ratio of foreground vs. background volume ($\theta_{\text{fg/bg, vol}} = (0.02, 0.048, \infty, \infty)^\top$) and the ratio of foreground vs. background intensity ($\theta_{\text{fg/bg, int}} = (1.0, 1.2, \infty, \infty)^\top$). Parameter values were identified using minimum, maximum, 5%-quantile and 95%-quantile values of the feature histograms as described in Sec. 4.3.1. The fuzzy sets to discard detections in poor quality image regions (E-F) were parameterized with $\theta_{r=0} = (-\infty, -\infty, c_z, c_z + 100)^\top$ and $\theta_{r=180} = (c_z - 100, c_z, \infty, \infty)^\top$, respectively, with c_z being the z -component of the embryo centroid. Thus, objects could be excluded based on their distance to the respective detection objective with a transition region in the center of the embryo. Note that fuzzy set parameters $-\infty$ or ∞ indicate plateau regions of the fuzzy sets towards the lower or the upper value range, respectively.

Additionally, depending on the rotation of the image, the membership values were reduced if a detected object was found in an imaged region of poor quality. This was accomplished by using the axial centroid location c_z of the embryo (z -component of the mean centroid location of all detected nuclei), to define an

additional fuzzy set for each of the rotations. Regions towards the respective detection objective obtained higher membership values, whereas low membership values were assigned for regions that were farther away (Fig. 6.5E, F). The individual membership functions of the features of every detected segment were combined using a fuzzy conjunction and potential false positive objects with a combined membership value of less than 0.5 were discarded. To remove redundant detections, objects that were detected in both views were combined using a weighted sum of the feature vectors if their distance was smaller than their minor axis [106]. The weighted average for the object combination was based on the fuzzy set membership degree (Sec. 2.4.3). Based on the artificial benchmark validation presented in Sec. 4.4 this approach proved to work well in practice (Tab. 4.7, TWANG+U+SF (SeMV), F-Score of 0.89). For the real microscopy data, however, the quality could only be assessed using visual inspection of an overlay of the fusion results and the maximum intensity projection due to the immense number of objects (up to ≈ 25000 cells for the investigated zebrafish embryos).

Tracking

Object correspondences in neighboring frames were identified using a nearest neighbor tracking approach as described in Sec. 4.5 and [106, 227, 244]. Depending on movement speed of the investigated nuclei, the maximum allowed distance for the nearest neighbor matching was set to 15-30 pixels, which corresponded to 0.1-0.2 $\mu\text{m}/s$ for the considered image material [106, 227]. To correct tracks where an object was missing in a single frame but re-appeared in the next frame, the missing object location was interpolated. However, to obtain tracks that ranged over a larger time span all available tracklets (short partial tracks of a single nucleus) were extracted and a subsequent nearest neighbor matching was employed to connect tracklet terminals both in the spatial and the temporal domain. The maximally allowed temporal displacement was set to 300 s and a maximum displacement speed of 0.17 $\mu\text{m}/s$ was used [106]. These parameters were identified via manual investigation of average cell movements in consecutive images and represent the maximum distance a cell can travel without causing ambiguous correspondences. Following the assumption that closely located nuclei migrate to the same destination, a branching of tracks was permitted if a maximum of two reasonable matching partners had been identified

at the terminal of a tracklet [227]. It should be noted, though, that these associations do not necessarily correspond to real mitotic events of the nuclei but helped to identify tracks that range over the entire time window and thus provided a convenient way to highlight regional migration patterns. A customized graphical user interface was used to qualitatively assess the correct operation of the automatic tracking by superimposing selected traces of individual nuclei on the respective maximum intensity projection images (Fig. 5.2).

6.2.3 Results

A comparison of the performance required by a bunch of nucleus segmentation methods on the described light-sheet microscopy images of zebrafish embryos was performed by analyzing differently sized crop regions of a full-size 3D image stack. The time needed to process 3D images with resolutions of $256 \times 256 \times 50$, $512 \times 512 \times 100$, $1024 \times 1024 \times 200$, $2048 \times 2048 \times 400$ voxels for sizes S, M, L and XL, respectively, was measured in seconds. The plain OTSU method was excluded from this test, as it did not produce satisfactory results in any of the benchmarks presented in Sec. 3.3 and [220]. Fig. 6.6 summarizes the performance obtained by OTSUWW, GAC, GFT, GC and TWANG to segment a single stack of the differently sized 3D image stacks. TWANG clearly dominated the field and was up to one order of magnitude faster than the compared algorithms. Even using a non-threaded implementation, the performance benefit of TWANG held true. Despite the good results of OTSUWW achieved on the 2D and 3D benchmark images [220], OTSUWW was 3-fold slower than the single threaded TWANG and up to 10-fold slower compared to our threaded implementation of TWANG applied on the L dataset due to the missing parallelization that becomes more significant for large images. Additionally, TWANG has a relatively low memory footprint due to the local extraction of objects. This property allowed even to process the largest image category (XL), which could not be handled by any of the compared algorithm within the memory limitation of 32GB. Analyzing a single experiment of about 10 TB of image data required approximately 12 hours on a small computing cluster consisting of 30 nodes (App. B). The same task would have taken multiple weeks up to months with previously existing approaches. As shown in Fig. 6.7, the obtained segmentation quality was comparable to the results of the more complex segmentation

algorithms. Due to the lack of a ground truth for the 3D+t data, only a qualitative comparison could be performed.




Figure 6.6: Processing times for the compared segmentation algorithms in seconds (lower values are better) applied on 3D microscopy images of an early zebrafish embryo with fluorescently labeled cell nuclei. Image sizes correspond to 256x256x50 (S), 512x512x100 (M), 1024x1024x200 (L) and 2048x2048x400 (XL) voxels. Missing bars indicate that the respective algorithms failed to process the given image size (sizes M, L and XL for Gradient Flow Tracking and only size XL for Geodesic Active Contours (*), Graph-Cuts and Otsu+Watershed (*)). The opposed direction of TWANG (8 Threads) for image size S originated from a processing time below one second. In all tested categories, TWANG segmentation delivered peak performance and was the only method that could process the XL images with the given memory constraint of 32GB (adapted from [220]).

The obtained segmentation quality of different investigated algorithms on an exemplary slice of a zebrafish dataset is shown in Fig. 6.7. Moreover, the results of the seed detection and the segmentation applied on the whole embryo level are provided in Fig. 6.8A, B. Both visualizations emphasize the suitability of the proposed TWANG segmentation for a reliable extraction of fluorescently labeled nuclei and the associated properties from large-scale 3D+t image data. An exemplary visualization of a selection of tracked nuclei is shown in Fig. 6.8C. In Fig. 6.8D, several snapshots of an early zebrafish embryo are visualized by plotting the azimuth and elevation angles derived from a spherical projection of all nuclei in a 2D plot [244]. This digitally performed unwrapping of the embryo allowed to map the 3D organism to a 2D representation. The color code used in Fig. 6.8D reflects the density of nuclei in their local environment, *i.e.*, blueish values indicate a low density, whereas reddish values indicate higher density. If combined with a tracking approach, this color code

can also be used to visualize developmental quantities such as migration speed, migration direction and division events of cells or even the cellular ancestry as demonstrated in [5, 216].




Figure 6.7: Segmentation quality achieved on a 3D image of labeled nuclei of a zebrafish embryo. The panels show the maximum intensity projections of 3 neighboring z-slices. (A) Original image, (B) adaptive thresholding using Otsu’s method [169], (C) Otsu’s method combined with watershed-based blob splitting [15, 169], (D) geodesic active contours [39], (E) gradient vector flow tracking [125], (F) graph-cuts segmentation [1], (G) TWANG segmentation and (H) a false-colored original image. Segmentation errors indicated by the symbols are split (#), merged (+), missing (o) or spurious (~) [220]. Scale bar: 10 μm .

6.2.4 Discussion

In this section, the entire pipeline for the automated analysis of large-scale light-sheet microscopy image data was presented. The newly developed TWANG method implemented in XPIWIT represents a reasonable approach to obtain a reliable and fast segmentation of labeled nuclei in distributed computing environments. Furthermore, a new information fusion approach was presented that is capable of combining information from different complementary views of a specimen. Extracted cell centroids were temporally connected to obtain detailed cell trajectories that allow characterizing cellular dynamics on the whole-

embryo level. Due to the limited space, however, only proof-of-principle results were presented in this section. The described pipeline was successfully applied to multiple projects in the field of developmental biology. In particular, the developed methods were used to:

- count fluorescently labeled nuclei in the spinal cord of laterally oriented zebrafish embryos using confocal microscopy data [221]
- track fluorescently labeled neural crest precursor cells to their final destination in the eye of zebrafish using light-sheet microscopy data [227]
- assess the impact of Bisphenol A (BPA) onto early embryonic development of zebrafish using light-sheet microscopy data [170]
- quantify number, density and movement behavior of cells in the *mondo* morphant using light-sheet microscopy data [244]
- to derive an ensemble averaged reference embryo that can be used to precisely unveil developmental differences of mutant and wild-type embryos using light-sheet microscopy data [106].

In all scenarios, the new approach allowed to qualitatively and quantitatively investigate developmental processes at the single cell level, which are for instance, cellular movements, density changes, cell proliferation rates, cell counts and epiboly stage to name but a few. The detailed biological description and results can be found in the respective publications [106, 170, 221, 227, 244].

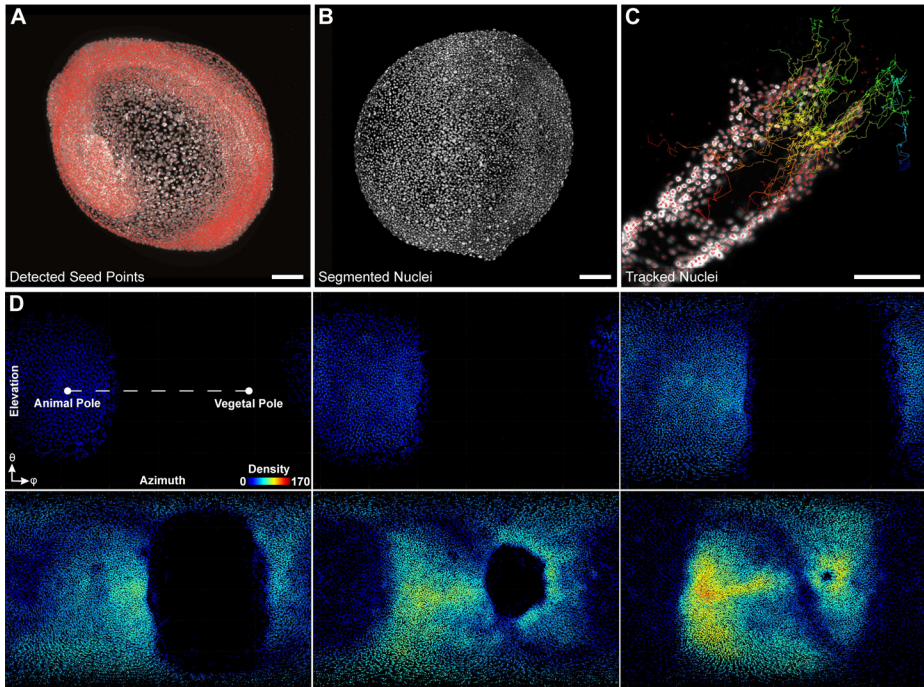


Figure 6.8: (A) Detected seed points using the LoG-based method described in Sec. 3.1 at 12 hpf. (B) Segmented nuclei using the TWANG segmentation algorithm as described in Sec. 3.3 at 11 hpf. (C) Traces of selected nuclei that were identified using a nearest neighbor tracking algorithm with tracklet fusion as described in Sec. 4.5. Centroids for the tracking were identified using the TWANG algorithm and the color code indicates time (increasing from red to blue). (D) Cell density distribution in a developing zebrafish embryo with a time difference of 200 minutes between each of the frames. Extracted coordinates of the labeled nuclei were transformed into spherical coordinates [244] and the panels show a selection of snapshots using the azimuth and elevation angles of each nucleus for plotting (time increases from left to right and top to bottom). Color encodes the number of neighbors of each nucleus that lie within a bounding sphere of 100px ($40\mu\text{m}$). Scale bar: $100\mu\text{m}$.

6.3 Automated Segmentation of Fluorescently Labeled Membrane Images

The third application example deals with the automatic segmentation of fluorescently labeled cellular membranes. Extracted cell shapes can significantly contribute to better understanding of cell-cell interactions, morphological changes during cell migration as well as phenotypic alterations on the whole embryo level [117, 167].

As indicated in Sec. 3.4, several methods for automatic membrane segmentation were presented in the past [63, 103, 151, 164]. However, all of the described methods rely on isotropic image data, *i.e.*, for anisotropic image data as mostly observed in fluorescence microscopy-based experiments, an upsampling of the data is performed to obtain isotropic resolution. Additionally, the use of computationally demanding processing operators renders most of the existing methods impractical for the analysis of large-scale experiments due to enormous memory consumption and processing times. In particular, terabyte-scale time series of light-sheet microscopy data that are routinely acquired in the field of developmental biology cannot be analyzed to their full extent and only excerpts of the data could be considered so far [103].

In this section, the newly developed RACE method for the segmentation of locally plane-like structures as described in Sec. 3.4 is applied to obtain cell shape information from large-scale images of fluorescently labeled membranes. The new algorithm is able to directly process anisotropic datasets, involves only efficient processing operators and offers a high degree of parallelization to perform new segmentations reliably and fast. At the same time, the segmentation quality is comparable or even better than the quality obtained by existing methods [216].

6.3.1 Dataset Description

The RACE algorithm was tested on different model organisms (fruit fly, zebrafish and mouse) and using both light-sheet (SiMView [230]) and confocal microscopy techniques (Zeiss 710 laser-scanning microscope). The respective image sizes of one channel of a single frame ranged from $640 \times 1300 \times 23$ (37 MB, 16 bit) for the confocal fruit fly dataset to $2048 \times 2048 \times 120$ (960 MB, 16 bit) for

a light-sheet image of a zebrafish embryo. All of the presented images contained both fluorescently labeled nuclei and membranes and were acquired at different developmental stages to test the algorithm under various different circumstances. Exemplary maximum intensity projections of the acquired image data are shown in Fig. 6.9A. Further information on the microscopy setup, the experimental procedure and the animal handling are provided in [216].

6.3.2 Automated Analysis Framework

The automatic membrane segmentation was performed using the new RACE algorithm as described in Sec. 3.4. Seeds were obtained from the membrane channel (MS) or from an additional nucleus channel (NS) and used for the fusion of the 2D segments that were identified by the slice-by-slice watershed algorithm applied on a membrane-enhanced input image. Essentially three intensity dependent parameters needed to be adapted and optimized for the different specimens and an overview of the manually identified optimal values are provided in App., Tab. D.3 and [216].

In addition to the performance validation on the simulated benchmark data presented in Sec. 3.4.3, a comparison to manually labeled ground truth images was performed on a set of representative image regions of different model organisms and using different microscopy techniques. The quantitative comparison confirmed the suitability of the proposed RACE approach with improved quality and speed compared to state-of-the-art methods. Further information about the ground truth generation, the comparison procedure as well as the details on the additional validation examples can be found in [216]. Due to the laboriousness of the manual ground truth generation, however, the algorithmic validation had to be performed on small representative image regions rather than on the whole embryo level.

6.3.3 Results

Volume renderings of the segmentation results produced by the RACE algorithm are shown in Fig. 6.9B, C, for different model organisms (fruit fly, zebrafish and mouse) and different microscopy techniques (light-sheet and confocal microscopy). In all investigated scenarios, RACE nicely extracted cell




Figure 6.9: Exemplary segmentation results obtained by applying the RACE algorithm to 3D light-sheet microscopy images for fruit fly, zebrafish and mouse embryos and a confocal image of a fruit fly embryo. (A) shows maximum intensity projections of each specimen. The remaining rows show volume renderings of the segmentation results from the outside of the embryo (B) and for a cross section that makes the interior of the embryos visible (C). The segmentation images show that the individual cells are nicely resolved, which allows to characterize cellular morphology for subsequent quantitative analyses. Scale bar: 50 μm (adapted from [216]).

shapes in the entire embryo that were suitable for further quantitative analyses. As demonstrated in Fig. 6.9C, even the interior of the hollow embryos only contained a negligible amount of false positive detections due to autofluorescence. All segmentation results could be processed within a few minutes on a single workstation computer instead of hours as required by existing methods. To quantitatively assess the time performance of the algorithms, differently sized crop regions of a fruit fly embryo were processed using the segmentation algorithms MARS [63], ACME [164], EDGE4D [103] and the different versions of the RACE algorithm [216]. The results of the performance comparison are depicted in Fig. 6.10. The RACE algorithm was by a large margin the fastest method, being one to two orders of magnitude faster than the compared algorithms. For a terabyte-scale dataset this essentially means that processing times are reduced from years to a few days with high segmentation accuracy.

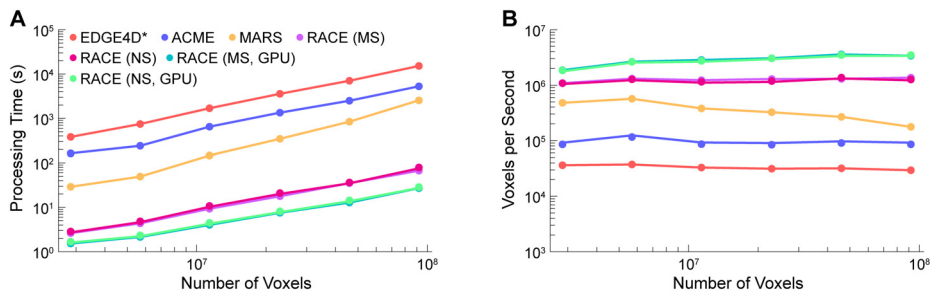


Figure 6.10: Performance comparison of the proposed RACE algorithm to other state-of-the-art membrane segmentation methods (ACME [164], MARS [63] and EDGE4D [103]). Differently sized crop regions of an early *Drosophila* embryo were processed and the obtained processing times in seconds are plotted against the respective number of voxels (A). In (B), the speed was measured by means of voxels per second. Note that ACME, MARS and EDGE4D work on isotropic image data, *i.e.*, the isotropic amount of voxels was used for the speed estimation. All versions of the RACE segmentation algorithm clearly outperformed the compared methods by at least an order of magnitude. The dedicated GPGPU implementation additionally pushes the performance by effectively eradicating existing bottlenecks of the algorithm. Note that the respective processing speed for each CPU-based and each GPU-based implementation of RACE was almost identical. Thus, the lines for RACE (NS) and RACE (MS) as well as RACE (NS,GPU) and RACE (MS, GPU) overlap in the performance plot and appear as a single line. Due to available implementation for Mac OS X, the performance of EDGE4D was assessed on an Apple MacBook Pro (Intel Core i7-3615QM CPU @ 2.3GHz, 16GB, Mac OS X 10.9.5) and is indicated by (*). For all other algorithms a workstation computer (2x Intel Xeon E5 CPUs @ 3.1 GHz, 196 GB RAM, NVidia Tesla K20 GPU, Windows 7 Professional 64-bit) was used to assess the performance (adapted from [216]).

In Fig. 6.11, an exemplary analysis of the obtained cell shape segmentation is shown. The visualization shows the color coded tissue anisotropy during *Drosophila* gastrulation, which can be nicely used to emphasize cell shape changes like the medio-lateral elongation of cells near the ventral furrow (Fig. 6.11A).

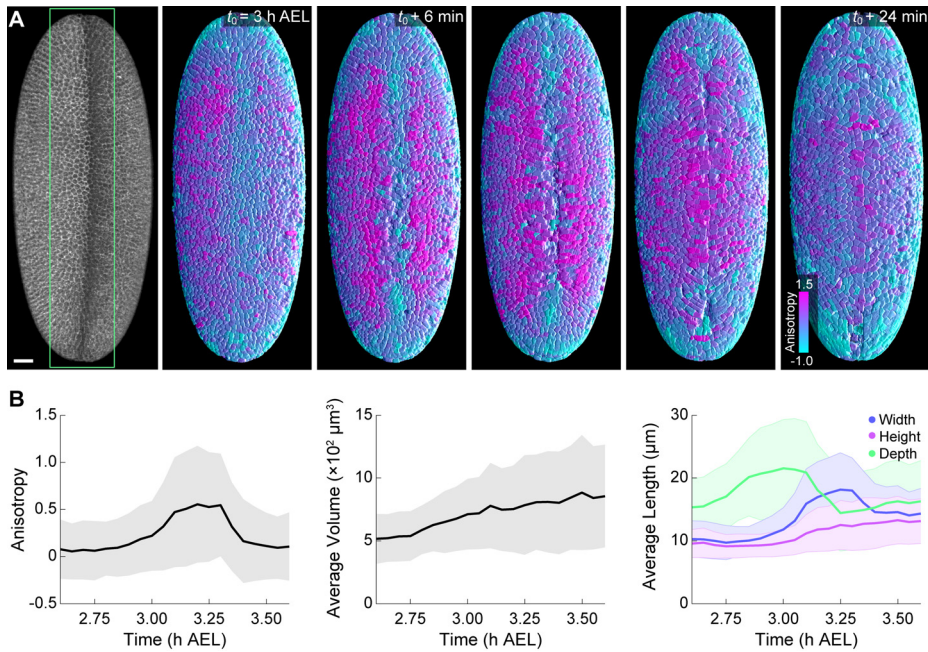


Figure 6.11: (A) Visualization of the tissue anisotropy observed during ventral fold formation of a *Drosophila* embryo starting at roughly 2.5 hours after egg laying (AEL). The color code reflects the anisotropy level of the cells, *i.e.*, cyan cells are elongated along the anteroposterior axis (height larger than width), purple cells are almost circular (similar height and width) and bright magenta cells are elongated along the mediolateral axis (width larger than height). In (B), average time courses of the anisotropy measure, the volume and the cell dimensions are shown for a selected region of interest (green rectangle in (A)). The furrow formation commences with a dorsoventral elongation of the cells (increased depth measure), followed by a mediolateral elongation of the cells (increased width measure). Upon closure of the ventral fold, both values decrease again. The same effect can be observed for the anisotropy measure, which serves as an indicator for mediolateral elongation of the cells. Shaded area indicates one standard deviation. Scale bar: $25 \mu\text{m}$ (adapted from [216]).

The anisotropy measure was calculated by

$$\text{Anisotropy} = \begin{cases} (w/h) - 1, & w/h > 1 \\ 1 - (h/w), & \text{otherwise} \end{cases}, \quad (6.1)$$

where w and h correspond to the medio-lateral and the anterior-posterior extent of the cells, respectively [216]. To obtain comparable extents at all locations of the embryo, each extracted cell was projected to an ellipsoid that was fit to the outer shell of the embryo. Extents were then estimated by projecting the segment to the medio-lateral tangent (width), the anterior-posterior tangent (height) and the surface normal (depth) at the respective centroid location. In addition to visualization of various features, properties of the extracted segments provided quantitative information of cell shape changes, such as anisotropy levels, volume or cell extents as shown in Fig. 6.11B. As demonstrated in [216], the approach could also successfully be linked to a nucleus tracking approach [5], which provides a combination of cellular shape information with the ancestry of each cell. In upcoming projects this could for instance be used to perform quantitative analysis of the temporal cell shape changes on the whole embryo level.

6.3.4 Discussion

In this section, exemplary results of the RACE segmentation method were presented. The method successfully segmented images of different specimens and different microscope types. While the results offered similar or even better segmentation quality, the new method represents a high-performance approach to cell shape segmentation with processing speed that exceeds the acquisition duration of common 3D microscopy images. In most cases where the fluorescently labeled nuclei were nicely distinguishable, the use of this additional channel for seed detection could directly improve the precision of the algorithm, by discarding many false positive detections in background regions [216]. However, the two channels had to be highly colocalized to guarantee that detected seeds were located within the correct compartment. For sequential acquisition of a membrane and a nucleus channel, for instance, it has to be ensured that

the acquired membrane image still encompasses the correct nucleus during the slightly later imaged second channel and vice versa.

In image regions with poor signal, the presented algorithm might potentially fail to close gaps in the object boundaries and produce 2D under-segmentation errors. Thus, more complex signal reconstruction approaches such as tensor voting might be inevitable if large gaps are present [151, 164]. An approach to potentially speed up the tensor voting operation would be to perform it also in a 2D slice-based manner as for the slice-by-slice watershed segmentation presented earlier. Moreover, the presented approach involves parameters related to the global image intensity. If the intensity of the input image heavily varies between slices it might be necessary to perform an additional preprocessing step that normalizes the intensity levels across the considered slices. An additional option might be to automatically determine separate intensity dependent threshold parameters for each of the slices.

In conclusion, the presented method is the first cell shape segmentation algorithm providing real-time processing capabilities to quantitatively analyze and visualize the vast amount of image data produced by modern 3D+t microscopy experiments in developmental biology.

7 Conclusions

As multidimensional image acquisition techniques are rapidly advancing, automated image analysis has become an indispensable tool to be able to handle the vast amount of generated image data. Especially in cases where datasets of a single experiment are comprised of multiple terabytes of image data, most existing algorithms tend to be far too computationally demanding or only offer limited result quality. Thus, the analyses had to be constrained to small subsets of the data or required an enormous amount of computing power to provide quantitative results reasonably fast. A central intention of the work summarized in this thesis was to develop new concepts and methods that permitted a reliable analysis of large-scale microscopy images significantly faster than existing methods without lowering the quality of automatically obtained results. The main idea to accomplish this goal was to rely on efficient processing operators where possible and to use as much available prior knowledge as possible to compensate potential weaknesses of the algorithms.

In Chap. 2, a general concept for the mathematical formulation of prior knowledge was proposed and it was shown, how image analysis pipelines can be equipped with the formalized prior knowledge to make more elaborate decisions. Furthermore, the framework includes the propagation of estimated result uncertainties, to be able to inform and improve downstream pipeline operators about the validity of their input data. Besides these general concepts, three new segmentation algorithms were developed that are able to reliably detect geometrical shapes like lines, spheres or plane-like structures in terabyte-scale multidimensional images (Chap. 3). To further improve the result quality, Chap. 4 demonstrated how existing pipelines can be systematically extended by uncertainty considerations, in order to filter, repair and fuse produced data. A quantitative assessment of the results produced by the extensions proofed their superior performance compared to the plain pipelines and thus render it a feasible approach to efficiently improve the result quality of image analysis pipelines. All methods were implemented in open-source software tools

(Chap. 5) and were successfully applied to analyze various challenging image-based experiments in the field of developmental biology (Chap. 6). In a nutshell, the major contributions presented in this thesis are:

1. A new concept on how to formalize, estimate and combine the uncertainty observed in a variety of image analysis problems based on fuzzy set membership degree (Chap. 2).
2. A general framework for the propagation of estimated uncertainty values in image analysis operators. This can be used to efficiently filter extracted data or to allow more complex information sharing compared to linear image analysis pipelines without uncertainty propagation (Sec. 2.4).
3. A highly efficient new algorithm for the detection and segmentation of (hyper-) spherical objects such as fluorescently labeled cell nuclei in large 3D microscopy images. The newly developed algorithm reaches segmentation quality of state-of-the-art algorithms while outperforming the existing approaches by a factor of up to ten with respect to processing time (Sec. 3.1, Sec. 3.3).
4. A new segmentation method that can extract comparable regions of interest along arbitrarily oriented line-like objects from multi-channel 2D images. A customized set of image features was developed in order to qualitatively and quantitatively compare intensity and object distributions within and between multiple color channels (Sec. 3.1, Sec. 3.2).
5. A new algorithm for the segmentation of locally plane-like structures such as fluorescently labeled cell membranes in 3D microscopy images. The algorithm was specifically developed for the application on large 3D datasets, *i.e.*, it was optimized to be memory efficient and fast compared to existing methods (Sec. 3.4).
6. The thorough validation of the functionality of the proposed methods on specialized validation benchmarks that were newly developed or adapted if suitable datasets were already available. The robustness of all methods was assessed using variable levels of image noise and image blur (Sec. 3.1-Sec. 3.4).

7. An extension of existing and newly developed image analysis operators to estimate and handle fuzzy uncertainty information. The quality of efficient operators could be significantly improved using the proposed fuzzy uncertainty propagation framework (Chap. 4).
8. A comprehensive new 3D cell simulation benchmark to model each step of an exemplary image analysis pipeline with the associated ground truth. This benchmark allows to draw conclusions on the performance of the entire processing pipeline instead of only considering individual operators independently (Sec. 4.1).
9. The implementation of all developed methods for the efficient analysis of 3D images in a new open-source software tool named XPIWIT that can be used for large-scale analyses of multidimensional images. The software runs on multiple platforms including Windows, Linux and Mac OS X, and can be executed on cluster systems like Hadoop. An easy-to-use XML-based pipeline definition format is used to execute, share and archive created image analysis protocols (Sec. 5.1).
10. Extensions of the open-source MATLAB toolbox Gait-CAD to be applicable for processing of multidimensional images (Sec. 5.2.1), object tracking (Sec. 5.2.2), quantification of neuron populations in the spinal cord of zebrafish (Sec. 5.2.3), trajectory visualization (Sec. 5.2.4), benchmark generation (Sec. 5.2.5) and semi-automatic parameter optimization (Sec. 5.2.6).
11. The successful application of the newly developed and implemented spinal cord extension for Gait-CAD to automatically analyze a high-throughput screen. In particular, the impact of more than 2000 small molecules on the neuronal development in the spinal cord of zebrafish embryos was quantitatively analyzed using the new software (Sec. 6.1).
12. An entire processing framework for 3D+t microscopy image data, which comprises registration, segmentation, tracking, visualization and statistical analysis. In close collaboration with the Institute of Toxicology and Genetics at KIT and Janelia Research Campus (Virginia, USA), hundreds of terabytes of 3D+t image data of fruit fly, zebrafish and mouse embryos were successfully analyzed and interpreted (Sec. 6.2, Sec. 6.3). In particular, features like cell counts, cell movements, cell shape changes as well

as density variations were analyzed in time-resolved 3D microscopy images. Thus, it was confirmed that the presented methods are capable of extracting quantitative information both reliably and fast even for large terabyte-scale datasets.

Although the image material considered throughout in this thesis came from different sources and was flawed with different acquisition deficiencies such as noise and blur, the disruptions of images within one dataset were comparable. In real experiments, however, the imaging conditions can vary significantly between different experiments or even during a single experiment. Causes for varying image content can be, for instance, temporally inhomogeneous expression of fluorescent reporters, varying experimental conditions or inherent variations between different probes or model organisms. Nevertheless, successful image analysis algorithms should be as robust as possible to be able to cope with these varying circumstances. An exemplary initiative aiming to provide standard methods for algorithmic validation and comparison is the Cell Tracking Challenge that was performed during the International Symposium on Biomedical Imaging 2013-2015 [144]. The provided datasets comprise simulated as well as manually annotated datasets of various specimens and microscope types to perform quantitative comparisons of cell segmentation and tracking algorithms. However, the covered simulated datasets tended to be too easy with a small amount of cells and real large-scale microscopy datasets, *e.g.*, of a fruit fly embryo could only be partially annotated due to the enormous manual effort required for the ground truth generation. Thus, more realistic simulated benchmark datasets are required in order to assess and compare the quality of different processing strategies in this challenging scenario. The artificial embryo simulation benchmark presented in this thesis and a recent extension of it presented in [217] already provide a first step in this direction but need to be further advanced to mimic even more realistic acquisition flaws at different levels of detail including tissue-dependent light scattering, realistic light attenuation as well as object distances and object sizes learned from real datasets.

To be able to cope with varying imaging conditions and different datasets, algorithmic parameters need to be optimized for the respective application scenario to obtain reasonable results. For the seed point detection step, the concept of interactive post-processing corrections was introduced in Sec. 5.2.6. Algorith-

mic parameters were initially adjusted such that false negative rates were minimized while the false positive rates potentially increased. In a post-correction step, the user could adapt various filter parameters to interactively clean the data. As the filtering was solely performed on the feature level, an instant visual feedback of the obtained results could be provided to the user without the necessity to reprocess any data on the image level. Thus, laborious parameter adjustments could be performed after the actual processing. This concept can potentially also be used to improve the result quality of other image analysis operators such as segmentation, multiview fusion and tracking, especially with respect to large-scale datasets of varying image content and quality. However, more sophisticated data structures are required, *e.g.*, to temporally store potential segmentation or tracking hypotheses that can be used for subsequent visualization and feature extraction. Moreover, a combination of filters can potentially be used to improve the discrimination capabilities and to be able to also resolve subtle object differences. The uncertainty framework presented Chap. 2 could be adapted to work with this kind of post-processing approach. In addition, an interesting extension would be to investigate if and how parameters of the fuzzy uncertainty framework can be automatically adapted or interpolated to match the specific dynamically changing requirements of the acquired image material. For instance, an interactive graphical user interface similar to the one presented in Sec. 5.2.6 could be used for sparse manual annotations that in turn help to automatically determine the uncertainty framework parameters before they are applied to the datasets.

With respect to the investigated tracking algorithms, there is still a need for even more accurate methods. Even the best published tracking algorithms have a finite error rate and are thus not yet capable of generating error-free cell lineages automatically without laborious manual proofreading and correction of identified correspondences [5, 61]. Hence, a substantial effort is required in the future to further improve the quality of object detection, object segmentation and the detection of events like cell divisions. The manual effort should be reduced to a minimum and supported with interactive and smart correction frameworks that guide the user to specific sites of potential errors and allow for easy corrections.

As microscopy systems are continuously advancing, it is important to regularly update developed implementations, *e.g.*, to make use of new hardware archi-

tectures, parallel programming paradigms or distributed computing solutions, to keep pace with latest developments. Especially, the recently presented time-resolved functional imaging techniques on the whole-embryo level will soon raise new challenges for fast and accurate automatic analysis. The membrane segmentation method presented in this thesis, for instance, provides real-time capabilities when executed on a high-end workstation. In the present state, however, the acquisition and the analysis were still considered as separate steps and were performed individually. An instant qualitative and quantitative feedback as well as live visualizations of biologically relevant features during the experiment would offer unprecedented possibilities to the experimentalists. For instance, to further improve the analysis of single-cell morphology changes, large-scale tissue reorganizations, temporal changes of gene expression patterns or signal molecule transduction at the whole-embryo level. This would allow to study effects of chemical treatments or other perturbations applied to the model systems in detail and at the time of the experiment. Being able to automatically extract error-free lineages and segmentations in real-time with a minimal amount or optimally no user intervention will ultimately enable truly interactive biological experiments on a daily basis.

A Nomenclature and Symbols

In this section the notation used throughout the thesis is summarized. The initial enumeration demonstrates the basic concepts used for nomenclature and the subsequent list of symbols explains all used symbols within their context. The nomenclature is partly based on the formulation given in [156].

- Lower case letters are used for scalars, parameters, indexing variables and functions (*e.g.*, a, b, c, d are common names for parameters, $f(\cdot), g(\cdot), h(\cdot)$ are common names for functions or functionals, i, j, k, l are common names for indices and finally x, y, z are common names for scalars or variables)
- Upper case letters are used for number of entities (*e.g.*, N, N_f for the number of data tuples and the number of features, respectively)
- Upper case bold face letters represent matrices (*e.g.*, \mathbf{M}, \mathbf{N})
- Lower case bold face letters represent vectors (*e.g.*, \mathbf{x}, \mathbf{y} for column vectors and $\mathbf{x}^\top, \mathbf{y}^\top$ for row vectors)
- Upper case script and Greek letters represent sets (*e.g.*, $\Omega, \Theta, \mathcal{X}, \mathcal{L}$)
- Accessing elements of data structures is denoted by square brackets (*e.g.*, $\mathbf{x}[1]$ accesses the first element of vector \mathbf{x} , $\mathbf{M}[1, 2]$ accesses the matrix element in the first row and second column)
- Elements topped with a hat refer to estimates of a certain quantity (*e.g.*, \hat{x} represents an estimate of the scalar x , $\hat{\mathbf{x}}$ represents an estimate of vector \mathbf{x})
- Elements topped with a tilde refer to a changed quantity (*e.g.*, $\tilde{\mathcal{X}}$ represents a modified subset of \mathcal{X})
- Optimal solutions are indicated by an asterisk (*e.g.*, x^* would denote the optimal solution to $\min(f(x))$)

Symbol	Description
$ x $	Absolute value of x
$\langle \mathbf{x}, \mathbf{y} \rangle$	Dot product of two vectors \mathbf{x} and \mathbf{y} , equivalent to $\mathbf{x}^\top \cdot \mathbf{y}$
*	The convolution operator. Discrete convolution is used for images.
$\ \mathbf{x} \ _2$	Euclidean norm of a vector \mathbf{x}
∇	The gradient operator. Finite differences are used to calculate image gradients.
\circ	Element-wise multiplication of two vectors or matrices (Hadamard product)
\wedge	Logical AND-operator
\vee	Logical OR-operator
$\frac{\partial f}{\partial x}$	Partial derivative of function f with respect to x
$\mathbf{0}$	Matrix or vector containing only zeros
$\mathbf{1}$	Matrix or vector containing only ones
2D	Two dimensional
2D+t	Two spatial dimensions and temporal dimension, e.g. 2D videos
3D	Three dimensional
3D+t	Three spatial dimensions and temporal dimension, e.g. 3D videos
AEL	After Egg Laying (used to specify the age of <i>Drosophila</i> embryos)
c	Index variable for image channels
c_z	z component of the centroid of an investigated specimen
CCD	Charge-coupled Devices
CMOS	Complementary Metal Oxide Semiconductor-based Detectors
CNR	Contrast-to-Noise Ratio
$\mathbf{d}_s(\mathbf{x})$	Vector pointing from seed location \mathbf{x}_s to \mathbf{x}
$D(i)$	Minimum distance of pixel i to the closest edge of a reference shape (used to calculate HM and NSD)
DNA	Deoxyribonucleic Acid

DoG	D ifference- o f- G aussian
DOM	D egree o f m embership in the context of fuzzy sets
DoM	D ifference- o f- M ean
DSL _M	D igital S canned L aser L ight S heet F luorescence M icroscopy
EDM	E uclidean D istance M ap
EDT	E uclidean D istance T ransform
$\exp(\cdot), e^x$	Exponential function
f	Index variable for features
FSMD	F uzzy S et M embership D egree
$\mathbf{g}_s(\mathbf{x}, \sigma_{\text{grad}})$	Image gradient for seed point s at location \mathbf{x} with regularization scale σ_{grad}
$\mathbf{G}(\mathbf{x}, \sigma)$	Gaussian kernel with standard deviation σ
g_i^*	Global intensity feature like min, max, mean, median or variance of channel c of an image
GFP	G reen F luorescent P rotein
GNU-GPL	G NU G eneral P ublic L icense
h	The height of maxima/minima that are suppressed by the h -maximum and h -minimum filters
H	2D histogram matrix containing label correspondences of a reference image and an automatic segmentation
HM	H ausdorff M etric
hpf	H ours P ost F ertilization (used to specify the age of zebrafish embryos)
i, j, k, l, m, n	Index variables defined by the context
ip_c	Area under the curve of an intensity profile p_c of channel c
ip_c^*	Features such as area under the curve, min, max, mean, median or variance of an intensity profile p_c of channel c
i_x, i_y, i_z	Index variables for x,y,z direction
I	General image variable
\mathbf{I}_r^{cl}	Iteratively closed image using a maximum structuring element radius r
\mathbf{I}^{cr}	Cropped image

\mathbf{I}_c^{cr}	Channel c of a cropped image
\mathbf{I}_s^{cr}	Cropped image for seed point s
\mathbf{I}^{dark}	Simulated dark current image of the detector
\mathbf{I}^{edm}	Euclidean distance map image
$\mathbf{I}^{\text{final}}$	Final simulated microscopy image containing all acquisition simulation steps
\mathbf{I}^{hee}	Image of the Hessian-based objectness filter response
\mathbf{I}^{iedm}	Euclidean distance map image of an inverted input image
\mathbf{I}^{LoG}	Laplacian-of-Gaussian filtered image
$\mathbf{I}^{\text{LoG}}(\mathbf{x}, \sigma)$	Laplacian-of-Gaussian with standard deviation σ at location \mathbf{x}
$\mathbf{I}^{\text{LoGMP}}$	Maximum intensity projection of multiple Laplacian-of-Gaussian filtered images
$\mathbf{I}^{\text{LoGMP}}(\mathbf{x}, \sigma_{\min}, \sigma_{\max})$	Laplacian-of-Gaussian maximum intensity projection of scales σ_{\min} to σ_{\max} at location \mathbf{x}
\mathbf{I}^{med}	Median filtered image defined by the context
\mathbf{I}^{MS}	Image containing the scale that yielded respective maximum value for $\mathbf{I}^{\text{LoGMP}}$
$\mathbf{I}^{\text{MS}}(\mathbf{x}, \sigma_{\min}, \sigma_{\max})$	The scale yielding the maximum value of $\text{LoGMP}(\mathbf{x}, \sigma_{\min}, \sigma_{\max})$
\mathbf{I}^{psf}	Image of a point spread function (PSF)
\mathbf{I}^{raw}	Unprocessed raw image
\mathbf{I}^{ref}	Reference image, <i>i.e.</i> , a labeled ground truth image
\mathbf{I}^{seg}	Segmentation image obtained with an automatic segmentation algorithm
\mathbf{I}^{ws}	Slice-based 2D watershed segmentation image
ITK	Insight Toolkit
Jl	Jaccard Index
k	Index variable for time-discrete time series
l	Index variable for linguistic terms
LoG	Laplacian-of-Gaussian
LSFM	Light-Sheet Fluorescence Microscopy
M	General number of data tuples, objects or dimensions specified by the context

MBF	Abbreviation for m embership f unction
$\min(\cdot), \max(\cdot)$	Minimum and maximum functions
n	Index variable for data tuples
$\mathbf{n}_s(\mathbf{x})$	Normal vector of seed point s at image location \mathbf{x}
N	General number of data tuples, objects or dimensions specified by the context
N_c	Number of channels of a multi-channel image
N_{cl}	General number of combined linguistic terms
$N_{cl,i}$	Number of combined linguistic terms for operator i
N_d	Number of dimensions
N_f	General number of features
$N_{f,i}$	Number of extracted features produced by operator i (i.e. columns in the result matrix)
N_i	Number of data tuples produced by operator i (i.e. rows in the result matrix)
N_k	Number of sampling points of time series or videos
N_l	General number of linguistic terms
$N_{l,i}$	Number of linguistic terms used by operator i
N_{ts}	Number of time series
$N_{ts,i}$	Number of time series produced by operator i
N_{op}	Number of operators within an analysis pipeline
N_p	Dimensionality of the a parameter vector, <i>i.e.</i> , the number of parameters
N_s	Number of seeds detected in an image
N_x, N_y, N_z	Number of pixels in x,y,z direction
$\mathcal{N}(\mu, \sigma)$	Normally distributed random variable with mean μ and standard deviation σ
\mathbf{nd}_c	Vector of variance normalized mean deviation of an intensity profile from a reference profile
NSD	N ormalized S um of D istances
$O(\lambda)$	Response of the objectness filter for the eigenvalue vector λ
\mathbf{p}_c	Vector containing the intensity profile of image channel c

$\bar{\mathbf{p}}_c^{\text{ctrl}}$	Vector containing the averaged intensity profiles of a control group of image channel c
PCL	P oint C loud L ibrary
PDE	P artial D ifferential E quation
plp_c	Peak location of intensity profile \mathbf{p}_c of channel c
PSF	P oint S pread F unction
pvp_c	Peak value of intensity profile \mathbf{p}_c of channel c
pwp_c	Peak width of intensity profile \mathbf{p}_c of channel c
\mathcal{P}_{bg}	The set of background pixels in a binary image
\mathcal{P}_{fg}	The set of foreground pixels in a binary image
$P_\lambda(\cdot)$	Poisson process to simulate photon shot noise using the image intensities as mean
q_l, q_u	Parameters for lower and upper quantile saturation
r	General variable for radius
\mathbb{R}	The set of real numbers
\mathbb{R}^{N_d}	The N_d -dimensional vector space of real numbers
r_{\min}, r_{\max}	Minimum and maximum radius of objects specified in the context
r_s	Approximate radius of seed point s detected in the LoGMP image
RACE	R ead-time A utomated C ell shape E xtractor
R_A, R_N, R_M	Parameters used for simulated object interaction terms
RFP	R ed F luorescent P rotein
RGB	R ed G reen B lue (color specification convention)
RI	R and I ndex
ROI	R egion of I nterest
$\mathbf{r}_x, \mathbf{r}_y$	Vectors containing the abscissa and the ordinate of a regression curve
s	Index variable for seed points
\mathbf{s}	Vector containing the physical spacing of an image along each dimension
S	Frobenius norm of the vector of eigenvalues
\mathbf{S}_r	Euclidean sphere structuring element with radius r
SDK	S oftware D evelopment K it
SNR	S ignal-to- N oise R atio

$\text{sgn}(\cdot)$	Signum function that returns the sign of a real number
SPIM	Selective/Single Plane Illumination Microscopy
t_{dbc}	Distance-based cutoff value used for feature fusion
t_{wmi}	Threshold for the window mean intensity feature
TWANG	Threshold of Weighted Intensity And Seed-Normal Gradient Dot Product Image
VTK	Visualization Toolkit
$w_s(\mathbf{x})$	TWANG weighting kernel at image location \mathbf{x} for seed point s
$w_{\text{adh}}, w_{\text{rep}}, w_{\text{bdr}}$	Weights for the adhesive displacement, the repulsive displacement and the boundary displacement
x	General variable defined by the context
\mathbf{x}	General vector variable defined by the context
\mathbf{x}_s	Position of seed point s
$x_i[n, l]$	l -th feature in the n -th data tuple produced by the i -th operator
$\mathbf{x}_i[n]$	n -th data tuple produced by the i -th operator containing $N_{f,i}$ feature entries
\mathbf{X}_i	Result matrix of the i -th operator containing N_i data tuples with $N_{f,i}$ features
\mathcal{X}_i	Variable for an output set of operator i
XML	Extensible Markup Language
XPIWIT	XML Pipeline Wrapper for the Insight Toolkit
α_{il}	Forward threshold of operator i and linguistic term l
β_{il}	Backwards threshold of operator i and linguistic term l
$\beta_{\text{hee}}, \gamma_{\text{hee}}$	Parameters to control the response of the Hessian-based objectness filter
$\Delta\mathbf{x}^{\text{adh}}$	Adhesive displacement vector of two neighboring objects
$\Delta\mathbf{x}^{\text{bdr}}$	Boundary constraint for dynamic object simulation
$\Delta\mathbf{x}^{\text{rep}}$	Repulsive displacement vector of two neighboring objects
$\Delta\mathbf{x}^{\text{tot}}$	Sum of all displacement vectors of a single object
θ	General parameter specified in the context
$\boldsymbol{\theta}$	General parameter vector specified in the context

$\lambda_1, \lambda_2, \dots, \lambda_{N_d}$	Eigenvalues sorted in ascending order of a $N_d \times N_d$ matrix defined in the context
λ	N_d -dimensional vector of eigenvalues sorted in ascending order
$\lambda(\mathbf{x}, \sigma_{\text{hee}})$	Vector of eigenvalues of the Hessian calculated at location \mathbf{x} of a σ_{hee} -regularized input image
μ	Arithmetic mean value
$\mu_{\mathcal{A}}(\cdot)$	Membership function for the fuzzy set \mathcal{A}
$\boldsymbol{\mu}_{\mathcal{A}_l}(\cdot)$	Vector of membership functions for all features of the l -th linguistic term
$\boldsymbol{\mu}_{\mathcal{A}_{il}}(\cdot)$	Vector of membership functions for all features of the l -th linguistic term of processing operator i
$\mu_{\mathcal{A}_{ifl}}(\cdot)$	Membership function of feature f and linguistic term l of the i -th processing operator
μ_c^{nd}	Average deviation of an intensity profile from an averaged reference profile
μ_{fg}	Mean intensity of the foreground signal (used for SNR calculation)
π	Mathematical constant (approximately 3.14159)
σ	Standard deviation
σ^2	Variance
σ_{agn}	Zero-mean additive Gaussian noise standard deviation
σ_{bg}	Standard deviation of the background signal of an image (used for SNR calculation)
σ_c^{ctrl}	Standard deviation of the control profiles of channel c
σ_{grad}	TWANG segmentation regularization scale used to smooth the vector field
σ_{hee}	Regularization scale of the Hessian-based objectness filter
σ_{kernel}	TWANG segmentation weighting kernel standard deviation
σ_{min}	Minimum scale used for the LoG-based seed detection
σ_{max}	Maximum scale used for the LoG-based seed detection
σ_{seed}	The scale of the LoG-filtered image where a seed point was detected at

σ_{smooth}	The standard deviation of a Gaussian smoothing kernel
σ_{smooth}^2	The variance of a Gaussian smoothing kernel
σ_{step}	Step size between adjacent scales used for the LoG-based seed detection
Σ	Covariance matrix
$\Sigma_{\text{row}}, \Sigma_{\text{col}}$	Vectors containing the row sums of all columns and the column sums of all rows of the 2D histogram \mathbf{H} , respectively
$\phi_s(\mathbf{x})$	Normalized dot product of image gradient and seed point normal at image location \mathbf{x} for seed point s
Φ	Variable for a set
$\chi_C(x)$	Characteristic function of the set C
ω_{kpm}	TWANG segmentation weighting kernel plateau multiplier
Ω_i	Variable for set that combines subsets of the outputs of operator i and operator $i - 1$
$\tilde{\Omega}_i$	Variable for a modified input set of operator $i + 2$

B Infrastructure

Although the algorithms and analysis methods presented in this thesis were developed with the focus on time and memory efficiency, the tremendous amount of multidimensional image data produced by modern microscopes demands for a specialized infrastructure to store, process, analyze and archive the acquired image data. This chapter introduces the platforms and services hosted by the Steinbuch Center for Computing (SCC) at the Karlsruhe Institute of Technology (KIT) that were used to establish the automated processing pipeline described in Chap. 6. Additionally, it is demonstrated how XPIWIT was used for parallel processing within this framework and which workarounds were necessary to cope with observed hardware limitations.

B.1 Large Scale Data Facility for Data Storage and Processing

The Large Scale Data Facility (LSDF) hosted at KIT has been providing storage, data management services as well as a computing facility that could be used for data analysis since 2009 [68]. It has been established in order to support large-scale scientific experiments, such as high-throughput microscopy as discussed in Chap. 6. Essentially, the LSDF framework provided storage capacity of multiple petabytes that was made accessible through high-speed network connections. Acquired microscopy data was transferred to this large storage system via optical fibers and was archived there for further processing and analysis of the images [106]. Data stored on the LSDF could be mounted locally via SSHFS (Secure SHell File System) or a Windows network drive to provide easy access to the raw data and the generated results.

To accelerate the result generation and automated analysis of the acquired image data, the processing was parallelized using an Apache Hadoop-based cluster for distributed computing hosted at SCC. The cluster consisted of two head

nodes for user interaction equipped with two quad-core Intel Xeon E5520 CPUs @ 2.27 GHz and 96 GB of main memory. The actual work was distributed over 30 worker nodes each containing the same processors as the head-nodes, 36 GB of memory, 2 TB hard disk drives and a Gigabit network connection. Scientific Linux 5.5 was used as an operating system on all nodes of the cluster. To be able to access raw image data and to store processed data, a gigabit network connection permitted fast data transfer between the LSDF and the respective processing nodes [67].

B.2 Hadoop Streaming to Parallelize XPIWIT

To speed up the image analysis pipeline implemented in XPIWIT, it was executed on a Hadoop-based cluster. Hadoop represents an open-source implementation of Google's map reduce framework and allows to easily parallelize applications, in case the computations are divisible into map tasks that independently compute intermediate results and reduce tasks that combine intermediate results to a final outcome [55].

Usually, Hadoop is used by specialized Java implementations of the map and reduce functions on the basis of the Hadoop libraries for Java. As XPIWIT was implemented in C++, Hadoop's streaming functionality was used instead of a native implementation in Java. Hadoop streaming essentially allows to specify arbitrary executables to be used as mappers or reducers. For the segmentation of thousands of images, processing each of the individual images was considered as an independent map task and could thus be processed in parallel on different nodes. A reducer was not used in this case as the segmentation did not need to combine information of individual images. However, the reducer might be used to combine obtained results, *e.g.*, to perform object tracking after all separately executed mappers have successfully processed the required object centroids. To inform the Hadoop job about the data to be processed, an XPIWIT input file had to be generated for each of the images. A special purpose MATLAB GUI was created to facilitate this task and is shown in App., Fig. B.1. A selected input folder on the LSDF was parsed for valid image files and for each of the images a unique XPIWIT input file was generated. An exemplary input file containing the required input, output and XML paths on the LSDF

is listed in App., Lst. B.1. Once all input files had been generated and copied to the Hadoop file system (HDFS), the actual streaming job was started. In App., Lst. B.2, the command used to start the Hadoop streaming batch process is shown.

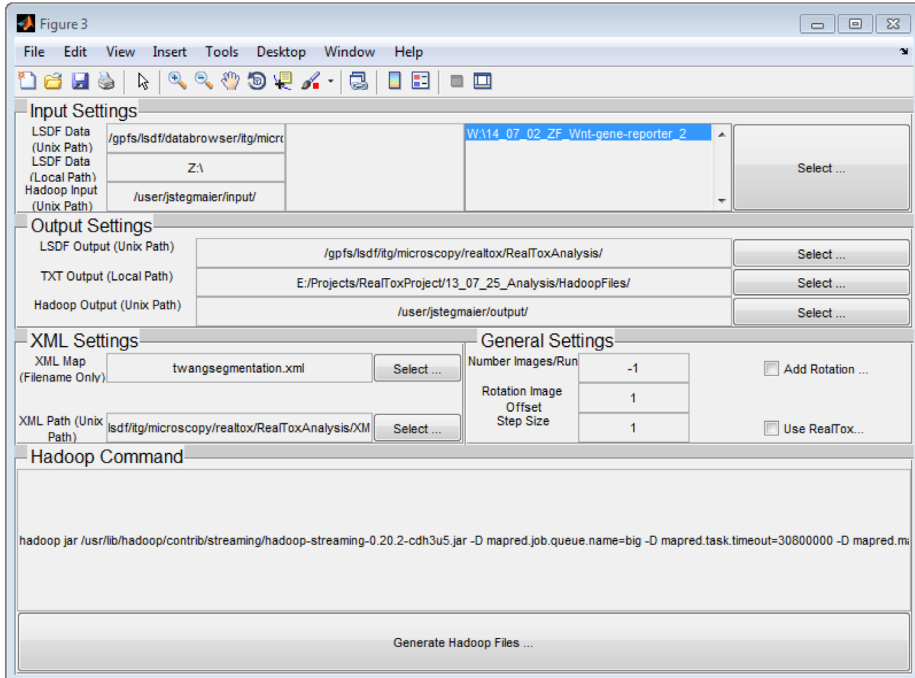


Figure B.1: Screenshot of the Hadoop input file generator implemented in MATLAB. Based on network drive access to the LSDF, the desired folders for input images, result output and the XML pipeline used for processing can be specified. Additionally, the Hadoop command and folder related settings can be adjusted if necessary.

Besides the essential call of the Hadoop streaming application, several parameters had to be provided to successfully start XPIWIT-based jobs on the cluster. Most importantly, the input directory, which was located on the HDFS, had to be specified. It contained the previously generated XPIWIT input files for each of the images about to be processed and a map task was generated for each of these files. The job scheduler took care of equally distributing the map tasks among all available processing nodes. The corresponding input parameters

for the mapper, reducer and their possible dependencies to shared libraries or other files are specified in Lines 9-14 in App., Lst. B.2.

```
1  --output /gpfs/lcdf/.../14_07_09_ZF_nc/Processed/,  
    result  
    --input 0, /gpfs/lcdf/.../14_07_09_ZF_nc/myimage_t=  
        0012.btf, 3, float  
3  --xml /gpfs/lcdf/.../XMLFiles/mysegmentation.xml  
    --seed 792  
5  --lockfile on  
    --subfolder filterid, filename  
7  --outputformat imagename, filename  
    --end
```

Listing B.1: XPIWIT Input Text Files Distributed by Hadoop.

Several problems were observed with the Hadoop streaming framework with respect to processing huge datasets. Most importantly, the job scheduler of the cluster was configured to start multiple map tasks at once on each of the cluster nodes. For applications with low memory demands, this might be a reasonable behavior, as the processing times generally should decrease by starting multiple jobs simultaneously. However, processing one of the large 3D images already required to make use of the entire main memory that was available on the respective cluster node. To prevent multiple tasks from being started concurrently on a single node, a so called lockfile-feature was added to XPIWIT, *i.e.*, a node specific text file was generated as soon as a node was occupied by an instance of XPIWIT (App., Lst. B.1, Line 5). Each idle process frequently checked if the lockfile for the node it was running on still was present on disk and remained in the idle state as long as the file existed. As soon as the lockfile got deleted, the next instance of XPIWIT started processing its data. However, one problem still occurred, if two or more processes wrote their lockfile exactly at the same time to the disk. In this case, both map-tasks were allowed to process and the node memory started swapping to disk. To prevent this rare event, a randomized sleep timer was used, which made the processes sleep for a short

random period of time, before they started to search for the lockfile. However, if processes were generated at the same time on the same node with exactly the same starting time used as seed for the random number generator, the problem persisted. Therefore, the random numbers were not seeded with the current time but with a special seed that was provided from the XPIWIT input file to each individual map-task (App., Lst. B.1, Line 4). This solved the issue of simultaneous processes on a single node and finally allowed to instantly start several thousands of tasks with a single call of Hadoop Streaming. Segmenting an entire embryo dataset with the TWANG algorithm described in Sec. 3.3 could be processed approximately within 12 hours using the described cluster and parallelization strategy. Performing the same task on a single high-end workstation computer would have required approximately two weeks, *i.e.*, the use of efficient parallelization is crucial in order to obtain experimental results in a contemporary way.

```
hadoop jar /usr/lib/.../hadoop/contrib/streaming/  
    hadoop-streaming-0.20.2-cdh3u5.jar  
2  -D mapred.map.max.attempts=50  
   -D mapred.job.queue.name=big  
4  -D mapred.task.timeout=30800000  
   -D mapred.map.tasks=1  
6  -D mapred.reduce.tasks=0  
   -input "/user/.../input/14_08_05_ZF_nc/"  
8  -output "/user/.../output/14_08_05_ZF_nc/"  
   -mapper "XPIWIT"  
10 -reducer 0  
   -file "/gpfs/lcdf/software/.../XPIWIT/XPIWIT"  
12 -file "/gpfs/lcdf/software/.../XPIWIT/libtiff.so.5"  
   -file "/gpfs/lcdf/software/.../XPIWIT/libmng.so.1"  
14 -file "/gpfs/lcdf/software/.../XPIWIT/liblcms.so.1"
```

Listing B.2: Hadoop Streaming Command used for Processing.

To further automate the image analysis and to facilitate the processing for non-experts, it is planned to directly integrate access to processing software via the data ingest tool developed at the Institute for Data Processing and Electronics (IPE) at KIT [90].

C Performance Assessment

An important component of algorithmic development is the quality assessment of the results produced by new methods and to compare these methods to existing approaches. The quantification and identification of errors allows to identify problems of algorithms and the possibility to improve their outcome by focusing on the particular weaknesses. This chapter summarizes the validation measures that were used to quantitatively compare seed detection, segmentation and tracking approaches as described in Chap. 3 and Chap. 4.

C.1 Seed Point Detection

To quantify the quality of a seed detection algorithm, the extracted spatial positions were compared to the original centroids of the ground truth objects. Due to the availability of segmentation ground truth, it was possible to simply check if a detected seed was localized within one of the object segments. If this was the case the seed was considered a *true positive (TP)* detection. If a detected seed was located in the background area, it was considered a *false positive (FP)*. The case of a *false negative (FN)* detection corresponds to an object present in the image data that was not associated to any of the detected seeds. The last remaining possibility, *i.e.*, a *true negative (TN)* is not considered, as in principal each background pixel that is not classified as a seed point corresponds to this case.

On the basis of the *TP*, *FP* and *FN*, the statistical measures recall, precision and F-Score were calculated as:

$$\text{recall} = \frac{TP}{TP + FN} \quad (\text{C.1})$$

$$\text{precision} = \frac{TP}{TP + FP} \quad (\text{C.2})$$

$$\text{F-Score} = 2 \cdot \frac{\text{recall} \cdot \text{precision}}{\text{recall} + \text{precision}} \quad (\text{C.3})$$

The values of recall, precision and F-Score all lie within the interval $[0, 1]$, with larger values being better. Recall reflects the relative amount of objects that were retrieved, *i.e.*, a value of 1 indicates that all objects were successfully identified. Contrary, precision measures the relative amount of correct objects out of all detected objects, *i.e.*, a value of 1 indicates that only correct objects and no false positives were detected. F-Score combines both measures to a single value and represents the harmonic mean of recall and precision.

For all correct detections, the Euclidean distance of the detected seed point to the centroid of the respective segment was additionally calculated to estimate how close the automatic detections are to the true object locations.

C.2 Segmentation

Contrary to the seed detection evaluation, the terms *false positive*, *false negative*, *true positive* and *true negative* are insufficient to solely quantify the quality obtained by a segmentation algorithm, *i.e.*, besides the topological errors objects might be correctly detected but possess different shapes or may be fragmented. To assess the correspondence between automatically derived segmentation results and the ground truth, the pixel-based Rand index (RI), the Jaccard index (JI), the Hausdorff metric (HM) and the normalized sum of distances (NSD) were used as defined in [50]. Additionally, topological errors are assessed as *split*, *merged*, *spurious* or *missing* objects. The extension of the 2D implementation by Coelho *et al.* [50] to 3D images is straightforward. The only necessary change is using a 3D implementation of the Euclidean distance map for distance calculations of segmented to reference pixels for HM and NSD.

Rand-Index (RI): Analogous to the definition in [50], \mathbf{I}^{ref} and \mathbf{I}^{seg} correspond to a labeled reference image and a labeled segmentation image, respectively,

with $\max(\mathbf{I}^{\text{ref}})$ and $\max(\mathbf{I}^{\text{seg}})$ unique labels for each of the detected objects. Moreover, let i, j be indexing all possible pixel pairs with $i \neq j$, then a, b, c, d are defined as the amount of pixel pairs that meet one of the criteria

- $(\mathbf{I}^{\text{ref}}[i] = \mathbf{I}^{\text{ref}}[j] \wedge \mathbf{I}^{\text{seg}}[i] = \mathbf{I}^{\text{seg}}[j]),$
- $(\mathbf{I}^{\text{ref}}[i] = \mathbf{I}^{\text{ref}}[j] \wedge \mathbf{I}^{\text{seg}}[i] \neq \mathbf{I}^{\text{seg}}[j]),$
- $(\mathbf{I}^{\text{ref}}[i] \neq \mathbf{I}^{\text{ref}}[j] \wedge \mathbf{I}^{\text{seg}}[i] = \mathbf{I}^{\text{seg}}[j]),$
- $(\mathbf{I}^{\text{ref}}[i] \neq \mathbf{I}^{\text{ref}}[j] \wedge \mathbf{I}^{\text{seg}}[i] \neq \mathbf{I}^{\text{seg}}[j]).$

Following the implementation of Coelho *et al.*, the values for a, b, c, d can be calculated using a 2D image histogram.

Considering for example the reference and segmentation images depicted in App., Fig. C.1, the label histogram listed in App., Tab. C.1 is obtained.

Histogram entries of the 2D histogram \mathbf{H} are generated by iterating over all pixel locations of the reference and the segmentation image and by using the respective object labels as column and row index. The histogram bins are increased by one for each index pair that points to them.

R/S	0	1	2	3	4	5	6	Σ_{row}
0	13922	0	0	0	0	0	24	13946
1	0	272	0	0	0	0	0	272
2	0	0	552	0	0	0	0	552
3	578	0	0	0	0	0	0	578
4	449	0	0	0	0	0	0	449
5	46	0	0	353	394	0	0	793
6	0	0	0	0	0	468	0	468
7	0	0	0	0	0	0	567	567
8	0	0	0	0	0	0	391	391
Σ_{col}	14995	272	552	353	394	468	982	

Table C.1: Histogram of the reference and segmentation images shown in App., Fig. C.1. Reference labels are determining the number of rows, whereas the segmentation image labels determine the number of columns. The label **0** is used as the image background label. Using the calculation scheme described in the main text yields $a = 97847186$, $b = 15489759$, $c = 507572$ and $d = 48434603$, which corresponds to a Rand Index of 0.90 (Eq. (C.8)) and a Jaccard Index of 2.27 (Eq. (C.9)).




Figure C.1: The image shows a labeled reference image (A) and an artificially distorted segmentation label image (B). In the segmentation image two nuclei were removed, two nuclei were merged and one nucleus was split. Numbers indicate the object IDs used as column and row indicators in the intersection histogram (App., Tab. C.1).

For lowest possible memory usage it is suggested to use sequentially increasing IDs for segments instead of uniformly distributed IDs over the whole value range, unless there really are that much objects to be evaluated. In the case of a valid ID range of 16 bit, this could for example lead to a histogram with 65535×65535 bins, which would require 16 GB of memory to be processed. To calculate the values a, b, c, d from the histogram, the binomial coefficient can be used. Thus, the first case can be calculated as

$$a = \sum_{i=0}^{\max(\mathbf{I}^{\text{seg}})} \sum_{j=0}^{\max(\mathbf{I}^{\text{ref}})} \binom{\mathbf{H}[i, j]}{2} \quad (\text{C.4})$$

for histogram entries $\mathbf{H}[i, j] > 1$. This reflects the sum of all possible index combinations that have equal labels in reference and segmentation, respectively. Similarly, the values for b and c can be calculated using the vectors for row and the column sums $\Sigma_{\text{row}}, \Sigma_{\text{col}}$:

$$b = \left(\sum_{i=0}^{\max(\mathbf{I}^{\text{ref}})} \binom{\Sigma_{\text{row}}[i]}{2} \right) - a \quad (\text{C.5})$$

$$c = \left(\sum_{j=0}^{\max(\mathbf{I}^{\text{seg}})} \binom{\Sigma_{\text{col}}[j]}{2} \right) - a \quad (\text{C.6})$$

The value b is the number of index pairs with $(\mathbf{I}^{\text{ref}}[i] = \mathbf{I}^{\text{ref}}[j] \wedge \mathbf{I}^{\text{seg}}[i] \neq \mathbf{I}^{\text{seg}}[j])$ and c is the number of index pairs with $(\mathbf{I}^{\text{ref}}[i] \neq \mathbf{I}^{\text{ref}}[j] \wedge \mathbf{I}^{\text{seg}}[i] = \mathbf{I}^{\text{seg}}[j])$. The value for d can be directly calculated from the remaining cases as:

$$d = \binom{N}{2} - a - b - c, \quad (\text{C.7})$$

with N being the number of pixels in the images. Finally, the Rand index [188] is defined by:

$$\text{RI}(\mathbf{I}^{\text{ref}}, \mathbf{I}^{\text{seg}}) = \frac{a + d}{a + b + c + d}, \quad (\text{C.8})$$

which is the fraction of index pairs that have the same labeling in reference and segmentation, respectively, versus all possible pixel pairs. This measure is normalized, *i.e.*, having a range of $[0, 1]$, where 1 indicates perfect agreement of all index pairs in the reference and the segmentation image [50]. In the result tables of this thesis, the RI values are scaled by 100, to obtain percentage values.

Jaccard-Index (JI): Using the same calculation of a, b, c, d as for the Rand index, the Jaccard index is defined as:

$$\text{JI}(\mathbf{I}^{\text{ref}}, \mathbf{I}^{\text{seg}}) = \frac{a + d}{b + c + d}. \quad (\text{C.9})$$

The higher the value attained by the Jaccard index, the better is the alignment of the reference image \mathbf{I}^{ref} to the segmentation image \mathbf{I}^{seg} [50].

Normalized Sum of Distances (NSD): The NSD is a normalized quantity that reflects the average distance of labeled pixels that do not agree in reference and segmentation:

$$\text{NSD}(\mathbf{I}^{\text{ref}}, \mathbf{I}^{\text{seg}}) = \frac{\sum_{i=1}^N ((\mathbf{I}^{\text{ref}}[i] \neq \mathbf{I}^{\text{seg}}[i]) \cdot D(i))}{\sum_{i=1}^N D(i)}. \quad (\text{C.10})$$

Here $D(i)$ is the minimum distance of pixel i to the edge of the reference shape being compared. The index i is running over all non-zero entries in the union

of images \mathbf{I}^{ref} and \mathbf{I}^{seg} . For the comparison ($\mathbf{I}^{\text{ref}}[i] \neq \mathbf{I}^{\text{seg}}[i]$) binary images are assumed. To avoid ambiguities of reference and segmentation, the calculation of the NSD is performed on individual labeled regions that are binarized and subsequently used for the NSD calculation. Values of the obtained NSD values of all considered objects are averaged for each image.

Hausdorff Metric (HM): To assess the deviation of reference and segmentation object, the Hausdorff metric is defined as the maximum of the set of minimal distances of two compared shapes [9].

$$\text{HM}(\mathbf{I}^{\text{ref}}, \mathbf{I}^{\text{seg}}) = \max_{i=1, \dots, N} \left(D(i) \cdot (\mathbf{I}^{\text{ref}}[i] \neq \mathbf{I}^{\text{seg}}[i]) \right). \quad (\text{C.11})$$

The measure is iteratively calculated for all overlapping shapes of reference and segmentation image and averaged by the number of reference objects afterwards. For both the NSD and the Hausdorff metric, the distance to the closest reference point at each pixel location i can be efficiently computed using the Euclidean distance map of the binarized reference image (\mathbf{I}^{edm}) and the inverted binarized reference image (\mathbf{I}^{iedm}) [53]:

$$D(i) = \max \begin{cases} \mathbf{I}^{\text{edm}}[i] & , \mathbf{I}^{\text{ref}}[i] \neq 0 \\ \mathbf{I}^{\text{iedm}}[i] & , \textit{otherwise} \end{cases} \quad (\text{C.12})$$

Two cases are required due to the fact that a misaligned segment can be either smaller or larger than the reference object. If the segmented object is smaller than the reference object, the maximum distance reflects the distance of the segmented object to the reference object from its inside ($\mathbf{I}^{\text{edm}}[i]$). In the other case, the distance has to be calculated from a point outside of the reference shape to its border ($\mathbf{I}^{\text{iedm}}[i]$).

The formulas for the evaluation measures are summarized in App., Tab. C.2.

Split Objects: Objects are considered as split if multiple fragments of the segmentation image are mapped to a single segment in the reference image. This can be again computed using the 2D histogram by searching for rows with multiple non-zero values (See *e.g.*, row 5 in C.1).

Measure	Formula	Range	Description
Rand Index (RI)	$\frac{a+d}{a+b+c+d}$	$[0, 1]$	Fraction of pixel combinations that have equal labels to all possible pixel combinations.
Jaccard Index (JI)	$\frac{a+d}{b+c+d}$	$[0, \infty)$	Fraction of pixel combinations that have equal labels to all non-equally labeled pixel combinations.
Hausdorff Metric (HM)	$\max (D(i) \cdot (\mathbf{I}^{\text{ref}}[i] \neq \mathbf{I}^{\text{seg}}[i]))$	$[0, \infty)$	Describes the maximum of the set of shortest distances between two segments [9].
Normalized Sum of Distances (NSD)	$\frac{\sum_{i=1}^N ((\mathbf{I}^{\text{ref}}[i] \neq \mathbf{I}^{\text{seg}}[i]) \cdot D(i))}{\sum_{i=1}^N D(i)}$	$[0, 1]$	The normalized distance of erroneously segmented pixels to the reference border [50].

Table C.2: Criteria for the quantitative comparison of a labeled ground truth segmentation with an automatically obtained segmentation image as defined in [9, 50, 188].

Merged Objects: Merged objects are defined as multiple objects in the reference image and as a single object in the segmentation image. By searching for columns with multiple non-zero values, these merged objects can be identified in the histogram representation (See *e.g.*, column 6 in C.1).

Spurious Objects: Objects that were erroneously added by the segmentation algorithm can be identified by searching for columns that have their maximum value at the background row of the reference.

Missing Objects: Missing objects can be identified in the 2D histogram by searching for rows, where the maximum values lie in the column of the segmentation background. In C.1, the third and fourth row correspond to missing objects of C.1.

Recall, Precision and F-Score: The values for recall, precision and F-Score as defined in Eq. (C.1), Eq. (C.2) and Eq. (C.3) were also calculated for the segmentation evaluation by considering correctly identified objects as true posi-

tives (TP), missing and merged objects as false negative detections (FN), and split and added objects as false positive detections (FP).

C.3 Tracking

To evaluate and compare the quality of the investigated tracking algorithms, the TRA measure described by Maška *et al.* was used [144]. This measure is calculated by considering the tracking result as an acyclic oriented graph. Each of the nodes within the graph corresponds to a detected object in a single frame and the edges between objects correspond to their temporal association. The idea of the tracking evaluation measure is to compare the graph of a manually annotated ground truth dataset to the automatically generated tracking graph. A high-quality automatic tracking should contain as few as possible differences to the ground truth graph, *i.e.*, counting and weighting the number of changes necessary to transform the automatically obtained graph to the ground truth graph can be used to assess the tracking quality. Due to the fact that some correction operations are more difficult than others, the required correction steps are weighted by the effort needed for manual correction. The weights defined in [144] are as follows with the required operation in brackets:

- delete a false positive node or edge: $w_{FP} = 1$ or $w_{ED} = 1$ (single deletion of a node (FP) or edge (ED))
- change an erroneous edge: $w_{EC} = 1$ (single reassignment of the edge terminal (EC))
- add a missing temporal association: $w_{EA} = 1.5$ (linkage of two objects necessary (EA))
- split an under-segmented node: $w_{NS} = 5$ (manual splitting the merged object (NS))
- add a missing object: $w_{FN} = 10$ (manual segmentation of the missing object (FN))

The weighted sum of required change operations can then be formulated as

$$\text{TRA}_P = w_{\text{NS}} \cdot \text{NS} + w_{\text{FN}} \cdot \text{FN} + w_{\text{FP}} \cdot \text{FP} + w_{\text{ED}} \cdot \text{ED} + w_{\text{EA}} \cdot \text{EA} + w_{\text{EC}} \cdot \text{EC}. \quad (\text{C.13})$$

The obtained TRA_P value is compared to the number of operations that would be required to manually generate the ground truth graph

$$\text{TRA}_E = w_{\text{FN}} \cdot |N| + w_{\text{EA}} \cdot |E|, \quad (\text{C.14})$$

with $|N|$ and $|E|$ being the number of nodes and edges, respectively. Finally, a normalization of the measure is performed using the following formulation:

$$\text{TRA} = 1 - \frac{\min(\text{TRA}_P, \text{TRA}_E)}{\text{TRA}_E}. \quad (\text{C.15})$$

The value range of the TRA is thus scaled to a range of $[0, 1]$ to be able to compare results obtained on different datasets with higher values being better and with 1 being ideal. In addition to this single value quality measure, a more detailed view on the errors is obtained by looking at the number of false positive detections, false negative detections, incorrect edges, missing edges, redundant edges and merged objects. This information is also provided by the respective log file produced for the TRA measure calculation. A more detailed description of the measure, implementation details and exemplary calculations are provided in the original publication [144].

C.4 Evaluation Platform

Unless otherwise stated, all measurements with respect to processing times were performed on a desktop PC equipped with an Intel Core i7-2600 CPU @ 3.4GHz and 32GB of memory installed using the Windows 7 (x64) operating system.

D Benchmark Datasets and Parameters

Name	Im.	N./B.	Resolution	σ_{agn}	σ_{smooth}^2	Description
SBDS1	30	1	314×220×50	-	-	3D images of simulated fluorescent nuclei by Svoboda <i>et al.</i> [226]. Used in Sec. 3.1 and Sec. 3.3.
SBDS2	50	10	314×220×50	[0.0, 0.128]	-	Selected images of SBDS1 with 10 different noise levels. Used in Sec. 3.1 and Sec. 3.3.
SBDS3	50	10	314×220×50	-	[0, 100]	Selected images of SBDS1 with 10 different blur levels. Used in Sec. 3.1 and Sec. 3.3.
SBDL	100	5	1024×768×2	[10 ⁻⁵ , 10 ⁻¹]	-	Simulated multi-channel 2D images containing elongated objects with spherical objects in the close vicinity. Used in Sec. 3.2.
SBDP1	4	1	128×128×51	-	-	3D images of simulated fluorescent membrane structures by Mosaliganti <i>et al.</i> [164]. Used in Sec. 3.4.
SBDP2	40	10	128×128×51	[0.0, 0.008]	-	Images of SBDP1 with 10 different noise levels. Used in Sec. 3.4.
SBDP3	40	10	128×128×51	-	[0, 100]	Images of SBDP1 with 10 different blur levels. Used in Sec. 3.4.

Table D.1: The benchmark datasets used for the validation experiments in Chap. 3. Columns list the dataset identifier (Name), the number of images (Im.), the number of noise/blur levels (N./B.), the image resolution, the noise/blur parameter ranges for σ_{agn} and σ_{smooth}^2 and the dataset description.

D Benchmark Datasets and Parameters

Name	Im.	N.	Objects	Resolution	σ_{agn}	Description
SBDE1	10	1	300–1000	640×640×128	0.0007	Five pairs of subsequent frames with different numbers of objects. Used for seed detection and segmentation validation (Sec. 4.2 and Sec. 4.3).
SBDE2	20	20	1000	640×640×128	[0.0005, 0.01]	Single frame of SBDE1 with 20 different noise levels. Used for seed detection and segmentation validation (Sec. 4.2 and Sec. 4.3).
SBDE3	10/20	1	300–1000	640×640×128	0.0007	Images of SBDS1 but with additional rotation simulation of sequential (SeMV) and simultaneous acquisition (SiMV). Images are rotated by 180° at every second frame for SeMV (10 images) and every image is rotated for SiMV (20 images). Used for the validation of the multi-view information fusion (Sec. 4.4).
SBDE4	100	1	1000	640×640×128	0.0007	Sequential time points containing 1000 objects each. Images are rotated by 180° at every frame for SiMV (2×50 images). Used for the validation of the tracking (Sec. 4.5).

Table D.2: The benchmark datasets used for the validation experiments in Chap. 4. Columns list the dataset identifier (Name), the number of images (Im.), the number of noise levels (N.), the number of objects (Objects), the image resolution, the noise parameter ranges for σ_{agn} and the dataset description.

Dataset	Drosophila (SIMView)		Mouse (SIMView)		Drosophila (Confocal)	
	RACE _{MS}	RACE _{NS}	RACE _{MS}	RACE _{NS}	RACE _{MS}	RACE _{NS}
Parameter/Algorithm						
XY vs. Z Ratio	5	5	5	5	5	5
2D Median Radius	2	2	2	2	2	2
Binary Threshold (MS)	0.0008	-	0.00004	-	0.002	-
Laplacian-of-Gaussian Sigma (NS)	-	4	-	4	-	4
Binary Threshold (NS)	-	OTSU	-	0.00002	-	OTSU
HMaximaFilter Height	20	20	40	45	20	20
HMaxima Binarization Threshold	0.00001	0.00001	0.00001	0.00001	0.00001	0.00001
HessianToObjectnessFilter Sigma	2	2	2	2	2	2
HessianToObjectnessFilter Alpha	1	1	1	1	1	1
HessianToObjectnessFilter Beta	1	1	1	1	1	1
HessianToObjectnessFilter Gamma	0.1	0.1	0.1	0.1	0.1	0.1
Morphological Closing MinRadius	4	4	1	1	4	4
Morphological Closing MaxRadius	4	4	4	4	4	4
Morphological Closing FilterMask3D	1	1	1	1	1	1
Slice-By-Slice Watershed Level	39	39	2	2	50	50
Slice-By-Slice RegionProps MinSeedArea	3	3	3	3	0	0
Slice-By-Slice FusionFilter MaxSegmentArea	1000	1000	2000	2000	1000	1000
Slice-By-Slice FusionFilter MinVolume	1500	1500	2300	2300	1500	1500
Slice-By-Slice FusionFilter MaxVolume	4000	4000	13000	13000	4000	4000

Table D.3: Parameterization of the RACE algorithm applied on fluorescence microscopy images of different specimens and microscopy techniques. The three most important intensity dependent parameters that need to be adjusted for each image are *Binary Threshold* (MS, NS), *HMaximaFilter Height* and *Slice-By-Slice Watershed Level*. The remaining parameters can be inferred from the experimental setup or may remain constant (adapted from [216]).

List of Figures

1.1	Fundamental principle of light-sheet fluorescence microscopy	4
1.2	Exemplary image analysis pipeline	8
1.3	Exemplary results of several image analysis methods	9
2.1	General image analysis pipeline concept	24
2.2	Different types of fuzzy set membership functions	29
2.3	Feature space partitioning using trapezoidal membership functions .	30
2.4	Visualization of fuzzy set membership degree for seed detection and segmentation	31
2.5	Extended image analysis pipeline concept with explicit uncertainty handling	37
3.1	General procedure for segmentation algorithm development	41
3.2	Benchmark Overview	42
3.3	Processing steps of a LoG-based seed detection	45
3.4	Processing steps of a LoG scale-space maximum projection for seed detection	45
3.5	Processing steps of an EDM-based seed detection	46
3.6	Euclidean distance map-based seed point detection	47
3.7	Assessment of the seed detection performance for different noise levels	50
3.8	Exemplary images of the SBDL benchmark dataset	53

3.9	Schematic overview of the elongated shapes segmentation	54
3.10	Processing steps and outcome of the elongated shape extraction . . .	56
3.11	Extraction quality of elongated objects on different noise levels	61
3.12	Schematic pipeline overview of the TWANG segmentation algorithm	64
3.13	Processing steps of the spherical object segmentation	66
3.14	TWANG segmentation weighting kernel	67
3.15	Quality assessment of the TWANG segmentation on a simulated 3D benchmark	71
3.16	Assessment of the 3D segmentation performance for different noise levels	72
3.17	Schematic pipeline overview of the RACE segmentation algorithm .	75
3.18	Benchmark Overview	76
3.19	Processing steps for a high-quality 2D segmentation	77
3.20	Euclidean distance map-based seed point detection	79
3.21	Processing steps for 3D segment fusion	81
3.22	Assessment of the RACE 3D segmentation performance for different noise levels	85
3.23	Qualitative comparison of the segmentation results of RACE, ACME and MARS	86
4.1	Exemplary image analysis pipeline	89
4.2	Pipeline schematic of the benchmark simulation	92
4.3	Maximum intensity projections of simulated benchmark images . . .	93
4.4	Potential functions of the object displacement	96
4.5	Visualization of the estimated uncertainty for the seed detection stage	103
4.6	Parameter sensitivity analysis for the different versions of a LoG- based seed detection	108

4.7	Assessment of the seed detection performance for different noise levels	109
4.8	Feature histograms and derived fuzzy set membership functions for segmentation classification	111
4.9	Uncertainty maps for the automatic segmentation	116
4.10	Assessment of the segmentation performance for different noise levels	118
4.11	Quantitative assessment of the tracking performance	129
5.1	Exemplary screenshot of the XPIWIT GUI on Windows and Ubuntu .	138
5.2	Exemplary screenshot of an Embryo3DT visualization	143
5.3	Exemplary screenshot of the semi-automatic data correction GUI . .	145
6.1	Overlay of RFP and GFP images of control embryos and LY-411575 treated embryos	149
6.2	Dose-response graphs for LY-411575 treated embryos	151
6.3	Light-sheet microscopy images of early zebrafish development . . .	156
6.4	Registration of beads detected in the Laplacian-of-Gaussian maximum projection	157
6.5	Membership functions for false positive reduction	159
6.6	Nucleus segmentation performance on differently sized 3D images .	162
6.7	Segmentation quality on a 3D image of labeled nuclei of a zebrafish embryo	163
6.8	Automatic analysis results for a developing zebrafish embryo	165
6.9	Whole embryo membrane segmentation results	168
6.10	Membrane segmentation performance comparison	169
6.11	Anisotropy maps of the ventral fold formation in <i>Drosophila</i> embryos	170
B.1	Screenshot of the Hadoop input file generator	191

C.1 Labeled reference image and an artificially distorted segmentation
image 198

List of Tables

1.1	Exemplary 3D+t datasets produced in embryomics	5
3.1	Quantitative assessment of the seed detection performance	49
3.2	Quantitative assessment of the detection quality for the regression curve estimation and the object detection quality	60
3.3	Comparison of the 3D Segmentation Quality for TWANG	70
3.4	Quantitative assessment of the segmentation performance	83
4.1	Benchmark parameters and generated ground truth data	94
4.2	Abbreviations, parameterizations and descriptions of the investigated seed detection algorithms	105
4.3	Quantitative assessment of the seed detection performance	106
4.4	Statistical quantities of the benchmark dataset	111
4.5	Abbreviations, parameters and descriptions of the segmentation algorithms	114
4.6	Quantitative assessment of the segmentation performance	117
4.7	Quantitative assessment of the multiview fusion performance	123
4.8	Quantitative assessment of the tracking performance	128
5.1	Comparison of image analysis software tools	139
6.1	Prior knowledge used for the algorithmic design	150
C.1	Intersection histogram for segmentation quality assessment	197

C.2	Performance evaluation criteria for segmentation algorithms	201
D.1	Benchmark datasets used for the validation experiments in Chap. 3	205
D.2	Benchmark datasets used for the validation experiments in Chap. 4	206
D.3	Parameterization of the RACE algorithm	207

List of Listings

5.1 Exemplary XML Pipeline.	135
B.1 XPIWIT Input Text Files Distributed by Hadoop.	192
B.2 Hadoop Streaming Command used for Processing.	193

Bibliography

- [1] AL-KOFAHI, Y.; LASSOUED, W.; LEE, W.; ROYSAM, B.: Improved Automatic Detection and Segmentation of Cell Nuclei in Histopathology Images. *IEEE Transactions on Biomedical Engineering* 57 (2010) 4, p. 841–852.
- [2] ALDRIDGE, B. B.; SAEZ-RODRIGUEZ, J.; MUHLICH, J. L.; SORGER, P. K.; LAUFFENBURGER, D. A.: Fuzzy Logic Analysis of Kinase Pathway Crosstalk in TNF/EGF/Insulin-induced Signaling. *PLoS Computational Biology* 5 (2009) 4, p. e1000340.
- [3] ALLAN, C.; BUREL, J.; MOORE, J.; BLACKBURN, C.; LINKERT, M.; LOYNTON, S.; MACDONALD, D.; MOORE, W.; NEVES, C.; PATTERSON, A.; PORTER, M.; TARKOWSKA, A.; LORANGER, B.; AVONDO, J.; LAGERSTEDT, I. AND. LIANAS, L.; LEO, S.; HANDS, K.; HAY, R. T.; PATWARDHAN, A.; BEST, C.; KLEYWEGT, G. J.; ZANETTI, G.; SWEDLOW, J. R.: OMERO: Flexible, Model-driven Data Management for Experimental Biology. *Nature Methods* 9 (2012) 3, p. 245–253.
- [4] AMAT, F.; KELLER, P. J.: Towards Comprehensive Cell Lineage Reconstructions in Complex Organisms using Light-Sheet Microscopy. *Development, Growth & Differentiation* 55 (2013) 4, p. 563–578.
- [5] AMAT, F.; LEMON, W.; MOSSING, D. P.; MCDOLE, K.; WAN, Y.; BRANSON, K.; MYERS, E. W.; KELLER, P. J.: Fast, Accurate Reconstruction of Cell Lineages from Large-Scale Fluorescence Microscopy Data. *Nature Methods* 11 (2014) 9, p. 951–958.
- [6] ANCHINI, R.; LIGUORI, C.; PAOLILLO, A.: Evaluation of the Uncertainty of Edge-Detector Algorithms. *IEEE Transactions on Instrumentation and Measurement* 56 (2007) 3, p. 681–688.
- [7] ANTIGA, L.: Generalizing Vesselness with Respect to Dimensionality and Shape. *The Insight Journal* 3 (2007), p. 1–14.

- [8] BALLARD, D.: Generalizing the Hough Transform to Detect Arbitrary Shapes. *Pattern Recognition* 13 (1981) 2, p. 111–122.
- [9] BAMFORD, P.: Empirical Comparison of Cell Segmentation Algorithms using an Annotated Dataset. In: *Proceedings of the International Conference on Image Processing*, Vol. 2, p. II–1073, IEEE, 2003.
- [10] BANDEMER, H.; GOTTWALD, S.: *Einführung in Fuzzy Methoden*. Akademie-Verlag, Berlin, 1993.
- [11] BAO, Z.; MURRAY, J.; BOYLE, T.; OOI, S.; SANDEL, M.; WATERSTON, R.: Automated Cell Lineage Tracing in *Caenorhabditis elegans*. *Proceedings of the National Academy of Sciences of the United States of America* 103 (2006) 8, p. 2707–2712.
- [12] BARTSCHAT, A.: *Analyse und Implementierung von digitalen Bildverarbeitungsfiltern zur Zellkerndetektion im Hinblick auf Qualität, Aufwand und Parallelisierbarkeit*. Bachelorarbeit, Duale Hochschule Baden-Württemberg, Karlsruher Institut für Technologie, 2013.
- [13] BARTSCHAT, A.; HÜBNER, E.; REISCHL, M.; MIKUT, R.; STEGMAIER, J.: XPIWIT - An XML Pipeline Wrapper for the Insight Toolkit. *Bioinformatics* 32 (2016) 2, p. 315–317.
- [14] BAUML, M.; BERNARDIN, K.; FISCHER, M.; EKENEL, H. K.; STIEFELHAGEN, R.: Multi-Pose Face Recognition for Person Retrieval in Camera Networks. In: *IEEE International Conference on Advanced Video and Signal Based Surveillance*, p. 441–447, IEEE, 2010.
- [15] BEARE, R.; LEHMANN, G.: The Watershed Transform in ITK-Discussion and New Developments. *The Insight Journal* (2006), p. 1–24.
- [16] BEDE, B.: *Mathematics of Fuzzy Sets and Fuzzy Logic*. Berlin Heidelberg: Springer, 2013.
- [17] BENTON, M. A.; AKAM, M.; PAVLOPOULOS, A.: Cell and Tissue Dynamics during *Tribolium* Embryogenesis Revealed by Versatile Fluorescence Labeling Approaches. *Development* 140 (2013) 15, p. 3210–3220.

-
- [18] BERTHOLD, M. R.; CEBRON, N.; DILL, F.; GABRIEL, T. R.; KÖTTER, T.; MEINL, T.; OHL, P.; SIEB, C.; THIEL, K.; WISWEDEL, B.: KNIME: The Konstanz Information Miner. In: *Data Analysis, Machine Learning and Applications*, p. 319–326, Springer, 2008.
- [19] BETTA, G.; CAPRIGLIONE, D.; CORVINO, M.; LIGUORI, C.; PAOLILLO, A.: Estimation of Influence Quantities in Face Recognition. In: *IEEE International Instrumentation and Measurement Technology Conference*, p. 963–968, IEEE, 2012.
- [20] BETTA, G.; CAPRIGLIONE, D.; LIGUORI, C.; PAOLILLO, A.: Uncertainty Evaluation in Face Recognition Algorithms. In: *IEEE International Instrumentation and Measurement Technology Conference*, p. 1–6, 2011.
- [21] BEUCHER, S.; MEYER, F.: The Morphological Approach to Segmentation: The Watershed Transformation. In: *Mathematical Morphology in Image Processing* (DOUGHERTY, E., Ed.), p. 433–481, 1993.
- [22] BIPM; IEC; IFCC; ISO; IUPAC; IUPAP; OIML: *Guide to the Expression of Uncertainty in Measurement*. 1993.
- [23] BISE, R.; YIN, Z.; KANADE, T.: Reliable Cell Tracking by Global Data Association. In: *Proc., IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, p. 1004–1010, 2011.
- [24] BISHOP, G.; WELCH, G.: An Introduction to the Kalman Filter. *Course Proceedings of SIGGRAPH 8* (2001), p. 27599–3175.
- [25] BLACKMAN, S. S.: Multiple Hypothesis Tracking for Multiple Target Tracking. *IEEE Aerospace and Electronic Systems Magazine* 19 (2004) 1, p. 5–18.
- [26] BOSKOVITZ, V.; GUTERMAN, H.: An Adaptive Neuro-fuzzy System for Automatic Image Segmentation and Edge Detection. *IEEE Transactions on Fuzzy Systems* 10 (2002) 2, p. 247–262.
- [27] BOUCHON-MEUNIER, B.; YAGER, R. R.; ZADEH, L. A.: *Information, Uncertainty and Fusion*. Boston: Kluwer Academic, 2000.

- [28] BOURGINE, P.; ČUNDERLÍK, R.; DRBLÍKOVÁ-STAŠOVÁ, O.; MIKULA, K.; REMEŠÍKOVÁ, M.; PEYRIÉRAS, N.; RIZZI, B.; SARTI, A.: 4D Embryogenesis Image Analysis Using PDE Methods of Image Processing. *Kybernetika* 46 (2010) 2, p. 226–259.
- [29] BOWDEN, R.; KAEWTRAKULPONG, P.: Towards Automated Wide Area Visual Surveillance: Tracking Objects between Spatially-Separated, Uncalibrated Views. In: *IEEE Proceedings on Vision, Image and Signal Processing*, Vol. 152, p. 213–223, 2005.
- [30] BRADSKI, G.: The OpenCV Library. *Dr. Dobb's Journal of Software Tools* (2000).
- [31] BRANSON, K.; ROBIE, A. A.; BENDER, J.; PERONA, P.; DICKINSON, M. H.: High-Throughput Ethomics in Large Groups of *Drosophila*. *Nature Methods* 6 (2009) 6, p. 451–457.
- [32] BREIMAN, L.: Random Forests. *Machine Learning* 45 (2001) 1, p. 5–32.
- [33] BROWN, M.; LOWE, D. G.: Automatic Panoramic Image Stitching using Invariant Features. *International Journal for Computer Vision* 74 (2007) 1, p. 59–73.
- [34] BURGER, W.; BURGE, M. J.: *Digital Image Processing: An Algorithmic Introduction using Java*. Springer, 2009.
- [35] CALLEJAS BEDREGAL, B.; DIMURO, G. P.; ROCHA COSTA, A.: Interval Fuzzy Rule-based Hand Gesture Recognition. In: *12th GAMM-IMACS International Symposium on Scientific Computing, Computer Arithmetic and Validated Numerics*, p. 12–12, IEEE, 2006.
- [36] CANNY, J.: A Computational Approach to Edge Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-8*, No. 6 (1986), p. 679–698.
- [37] CARPENTER, A.; JONES, T.; LAMPRECHT, M.; CLARKE, C.; KANG, I.; FRIMAN, O.; GUERTIN, D.; CHANG, J.; LINDQUIST, R.; MOFFAT, J.; GOLLAND, P.; SABATINI, D. M.: CellProfiler: Image Analysis Software for Identifying and Quantifying Cell Phenotypes. *Genome Biology* 7 (2006) 10, p. R100.

- [38] CARRANZA, N.; SMAL, I.; DZYUBACHYK, O.; NIESSEN, W.; MEIJERING, E.: Automated Lineage Tree Reconstruction from *Caenorhabditis Elegans* Image Data using Particle Filtering based Cell Tracking. In: *Proc., IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, p. 1921–1924, 2011.
- [39] CASELLES, V.; KIMMEL, R.; SAPIRO, G.: Geodesic Active Contours. *International Journal of Computer Vision* 22 (1997) 1, p. 61–79.
- [40] CEDILNIK, A.; GEVECI, B.; MORELAND, K.; AHRENS, J.; FAVRE, J.: Remote Large Data Visualization in the ParaView Framework. In: *Proceedings of the 6th Eurographics conference on Parallel Graphics and Visualization*, p. 163–170, Eurographics Association, 2006.
- [41] DE CHAUMONT, F.; DALLONGEVILLE, S.; CHENOUEARD, N.; HERVÉ, N.; POP, S.; PROVOOST, T.; MEAS-YEDID, V.; PANKAJAKSHAN, P.; LECOMTE, T.; LE MONTAGNER, Y.; LAGACHE, T.; DUFOUR, A.; OLIVOMARIN, J. C.: Icy: An Open Bioimage Informatics Platform for Extended Reproducible Research. *Nature Methods* 9 (2012) 7, p. 690–696.
- [42] CHEN, J.; ZHAO, L.: A New Method for Uncertainty Evaluation of Centroid Detection. In: *2nd International Congress on Image and Signal Processing, 2009. CISP'09*, p. 1–4, IEEE, 2009.
- [43] CHEN, J.; ZHAO, L.: A New Method for Uncertainty Evaluation of Corner Detection. In: *International Conference on Artificial Intelligence and Computational Intelligence*, Vol. 1, p. 458–462, IEEE, 2010.
- [44] CHEN, S.; ZHU, Y.; XIA, W.; XIA, S.; XU, X.: Automated Analysis of Zebrafish Images for Phenotypic Changes in Drug Discovery. *Journal of Neuroscience Methods* 200 (2011) 2, p. 229–236.
- [45] CHENOUEARD, N.; SMAL, I.; DE CHAUMONT, F.; MAŠKA, M.; SBALZARINI, I. F.; GONG, Y.; CARDINALE, J.; CARTEL, C.; CORALUPPI, S.; WINTER, M.; ET AL.: Objective Comparison of Particle Tracking Methods. *Nature Methods* 11 (2014), p. 281–289.
- [46] CHHETRI, R. K.; AMAT, F.; WAN, Y.; HÖCKENDORF, B.; LEMON, W. C.; KELLER, P. J.: Whole-Animal Functional and Developmental Imaging

- with Isotropic Spatial Resolution. *Nature Methods* 12 (2015) 12, p. 1171–1178.
- [47] CHKLOVSKII, D. B.; VITALADEVUNI, S.; SCHEFFER, L. K.: Semi-Automated Reconstruction of Neural Circuits using Electron Microscopy. *Current Opinion in Neurobiology* 20 (2010) 5, p. 667–675.
- [48] CIESIELSKI, K. C.; MIRANDA, P.; FALCÃO, A.; UDUPA, J. K.: Joint Graph Cut and Relative Fuzzy Connectedness Image Segmentation Algorithm. *Medical Image Analysis* (2013).
- [49] CIRESAN, D.; MEIER, U.; SCHMIDHUBER, J.: Multi-Column Deep Neural Networks for Image Classification. In: *IEEE Conference on Computer Vision and Pattern Recognition*, p. 3642–3649, IEEE, 2012.
- [50] COELHO, L.; SHARIFF, A.; MURPHY, R.: Nuclear Segmentation in Microscope Cell Images: A Hand-Segmented Dataset and Comparison of Algorithms. In: *Proc., IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, p. 518–521, 2009.
- [51] COELHO, L. P.; GLORY-AFSHAR, E.; KANGAS, J.; QUINN, S.; SHARIFF, A.; MURPHY, R. F.: Principles of Bioimage Informatics: Focus on Machine Learning of Cell Patterns. In: *Linking Literature, Information, and Knowledge for Biology*, p. 8–18, Springer, 2010.
- [52] COLLINS, D. L.; HOLMES, C. J.; PETERS, T. M.; EVANS, A. C.: Automatic 3D Model-based Neuroanatomical Segmentation. *Human Brain Mapping* 3 (1995) 3, p. 190–208.
- [53] CUISENAIRE, O.; MACQ, B.: Fast Euclidean Distance Transformation by Propagation Using Multiple Neighborhoods. *Computer Vision and Image Understanding* 76 (1999) 2, p. 163–172.
- [54] DAS, P.; VEKSLER, O.; ZAVADSKY, V.; BOYKOV, Y.: Semiautomatic Segmentation with Compact Shape Prior. *Image and Vision Computing* 27 (2009) 1, p. 206–219.
- [55] DEAN, J.; GHEMAWAT, S.: MapReduce: Simplified Data Processing on Large Clusters. *Communications of the ACM* 51 (2008) 1, p. 107–113.

- [56] DERICHE, R.: Fast Algorithms for Low-Level Vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12 (1990) 1, p. 78–87.
- [57] DUBOIS, D.; PRADE, H. (Ed.): *Fundamentals of Fuzzy Sets*. Boston: Kluwer Academic, 2000.
- [58] EISSING, N.; HEGER, L.; BARANSKA, A.; CESNJEVAR, R.; BÜTTNER-HEROLD, M.; SÖDER, S.; HARTMANN, A.; HEIDKAMP, G. F.; DUDZIAK, D.: Easy Performance of 6-Color Confocal Immunofluorescence with 4-Laser Line Microscopes. *Immunology Letters* 161 (2014) 1, p. 1–5.
- [59] ELICEIRI, K.; BERTHOLD, M.; GOLDBERG, I.; IBÁÑEZ, L.; MANJUNATH, B.; MARTONE, M.; MURPHY, R.; PENG, H.; PLANT, A.; ROYSAM, B.; STUURMAN, N.; SWEDLOW, J. R.; PAVEL TOMANCAK, P.; CARPENTER, A. E.: Biological Imaging Software Tools. *Nature Methods* 9 (2012) 7, p. 697–710.
- [60] FAUQ, A. H.; SIMPSON, K.; MAHARVI, G. M.; GOLDE, T.; DAS, P.: A Multigram Chemical Synthesis of the γ -Secretase Inhibitor LY411575 and its Diastereoisomers. *Bioorganic & Medicinal Chemistry Letters* 17 (2007) 22, p. 6392–6395.
- [61] FAURE, E.; SAVY, T.; RIZZI, B.; MELANI, C.; STAŠOVÁ, O.; FABRÈGES, D.; ŠPIR, R.; HAMMONS, M.; ČÚNDERLÍK, R.; RECHER, G.; ET AL.: A Workflow to Process 3D+Time Microscopy Images of Developing Organisms and Reconstruct Their Cell Lineage. *Nature Communications* 7 (2016) 8674, p. 1–10.
- [62] FENT, K.; WEISBROD, C. J.; WIRTH-HELLER, A.; PIELES, U.: Assessment of Uptake and Toxicity of Fluorescent Silica Nanoparticles in Zebrafish (*Danio rerio*) Early Life Stages. *Aquatic Toxicology* 100 (2010) 2, p. 218–228.
- [63] FERNANDEZ, R.; DAS, P.; MIRABET, V.; MOSCARDI, E.; TRAAS, J.; VERDEIL, J.-L.; MALANDAIN, G.; GODIN, C.: Imaging Plant Growth in 4D: Robust Tissue Reconstruction and Lineaging at Cell Resolution. *Nature Methods* 7 (2010) 7, p. 547–553.
- [64] FRANGI, A. F.; NIESSEN, W. J.; VINCKEN, K. L.; VIERGEVER, M. A.: Multiscale Vessel Enhancement Filtering. In: *Medical Image Computing and Computer-Assisted Intervention*, p. 130–137, Springer, 1998.

- [65] FRANK, E.; HALL, M.; HOLMES, G.; KIRKBY, R.; PFAHRINGER, B.; WITTEN, I.: Weka: A Machine Learning Workbench for Data Mining. In: *Data Mining and Knowledge Discovery Handbook: A Complete Guide for Practitioners and Researchers*, p. 1305–1314, New York: Springer, 2005.
- [66] FREEDMAN, D.; ZHANG, T.: Interactive Graph Cut based Segmentation with Shape Priors. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Vol. 1, p. 755–762, IEEE, 2005.
- [67] GARCIA, A.; BOUROV, S.; HAMMAD, A.; JEJKAL, T.; OTTE, J.; PFEIFFER, S.; SCHENKER, T.; SCHMIDT, C.; VAN WEZEL, J.; NEUMAIR, B.; STREIT, A.: Data Management and Analysis at the Large Scale Data Facility. In: *Proc., 6th International Conference on Digital Information Management, Melbourne, 2011*.
- [68] GARCIA, A. O.; BOUROV, S.; HAMMAD, A.; VAN WEZEL, J.; NEUMAIR, B.; STREIT, A.; HARTMANN, V.; JEJKAL, T.; NEUBERGER, P.; STOTZKA, R.: The Large Scale Data Facility: Data Intensive Computing for Scientific Experiments. In: *Proc., IEEE International Symposium on Parallel and Distributed Processing Workshops and Phd Forum*, p. 1467–1474, IEEE, 2011.
- [69] GEHRIG, J.; REISCHL, M.; KALMAR, E.; FERG, M.; HADZHIEV, Y.; ZACKER, A.; SONG, C.; SCHINDLER, S.; LIEBEL, U.; MÜLLER, F.: Automated High Throughput Mapping of Promoter-Enhancer Interactions in Zebrafish Embryos. *Nature Methods* 6 (2009) 12, p. 911–916.
- [70] GELASCA, E. D.; OBARA, B.; FEDOROV, D.; KVILEKVAL, K.; MANJUNATH, B.: A Biosegmentation Benchmark for Evaluation of Bioimage Analysis Methods. *BMC Bioinformatics* 10 (2009) 1, p. 368.
- [71] GENOVESIO, A.; LIEDL, T.; EMILIANI, V.; PARAK, W.; COPPEY-MOISAN, M.; OLIVO-MARIN, J.: Multiple Particle Tracking in 3D+ t Microscopy: Method and Application to the Tracking of Endocytosed Quantum Dots. *IEEE Transactions on Image Processing* 15 (2006) 5, p. 1062–1070.
- [72] GHAYE, J.; DE MICHELI, G.; CARRARA, S.: Simulated Biological Cells for Receptor Counting in Fluorescence Imaging. *BioNanoScience* 2 (2012) 2, p. 94–103.

-
- [73] GODINEZ, W.; LAMPE, M.; EILS, R.; MULLER, B.; ROHR, K.: Tracking Multiple Particles in Fluorescence Microscopy Images via Probabilistic Data Association. In: *Proc., IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, p. 1925–1928, 2011.
- [74] GODINEZ, W.; LAMPE, M.; WÖRZ, S.; MÜLLER, B.; EILS, R.; ROHR, K.: Deterministic and Probabilistic Approaches for Tracking Virus Particles in Time-Lapse Fluorescence Microscopy Image Sequences. *Medical Image Analysis* 13 (2009) 2, p. 325–342.
- [75] GONZALEZ, R.; WOODS, R.; EDDINS, S.: *Digital Image Processing Using MATLAB*. Prentice-Hall, Inc. Upper Saddle River, NJ, USA, 2003.
- [76] GONZALEZ, R. C.; WOODS, R. E.: *Digital Image Processing*. Prentice Hall, 2007.
- [77] GRABNER, M.; GRABNER, H.; BISCHOF, H.: Fast Approximated SIFT. *Asian Conference on Computer Vision* (2006), p. 918–927.
- [78] GUALDA, E. J.; VALE, T.; ALMADA, P.; FEIJÓ, J. A.; MARTINS, G. G.; MORENO, N.: OpenSpinMicroscopy: An Open-Source Integrated Microscopy Platform. *Nature Methods* 10 (2013) 7, p. 599–600.
- [79] GUL-MOHAMMED, J.; ARGANDA-CARRERAS, I.; ANDREY, P.; GALY, V.; BOUDIER, T.: A Generic Classification-based Method for Segmentation of Nuclei in 3D Images of Early Embryos. *BMC Bioinformatics* 15 (2014) 1, p. 9.
- [80] HENDERSON, A.: *ParaView Guide, A Parallel Visualization Application*. Kitware Inc., 2007.
- [81] HILL, D. L.; BATCHELOR, P. G.; HOLDEN, M.; HAWKES, D. J.: Medical Image Registration. *Physics in Medicine and Biology* 46 (2001) 3, p. R1–R45.
- [82] HU, W.; XIE, D.; FU, Z.; ZENG, W.; MAYBANK, S.: Semantic-Based Surveillance Video Retrieval. *IEEE Transactions on Image Processing* 16 (2007) 4, p. 1168–1181.
- [83] HUANG, D.; PAN, W.: Incorporating Biological Knowledge into Distance-based Clustering Analysis of Microarray Gene Expression Data. *Bioinformatics* 22 (2006) 10, p. 1259–1268.

- [84] HÜBNER, E.: Implementation and Development of a Graphical User Interface for Increased Accessibility of Image Processing Pipelines. Tech. Rep., Karlsruhe Institute of Technology, 2014.
- [85] HUISKEN, J.; STAINIER, D.: Selective Plane Illumination Microscopy Techniques in Developmental Biology. *Development* 136 (2009) 12, p. 1963–1975.
- [86] HUISKEN, J.; SWOGER, J.; DEL BENE, F.; WITTBRODT, J.; STELZER, E.: Optical Sectioning Deep Inside Live Embryos by Selective Plane Illumination Microscopy. *Science* 305 (2004) 5686, p. 1007–1009.
- [87] IBANEZ, L.; SCHROEDER, W.; NG, L.; CATES, J.; ET AL.: *The ITK Software Guide*. Citeseer, 2005.
- [88] JÄHNE, B.: *Digitale Bildverarbeitung*. Springer-Verlag, 2012.
- [89] JAIN, V.; SEUNG, H. S.; TURAGA, S. C.: Machines that Learn to Segment Images: A Crucial Technology for Connectomics. *Current Opinion in Neurobiology* 20 (2010) 5, p. 653–666.
- [90] JEJKAL, T.; HARTMANN, V.; STOTZKA, R.; OTTE, J.; GARCIA, A.; VAN WEZEL, J.; STREIT, A.: LAMBDA—The LSDF Execution Framework for Data Intensive Applications. In: *Proc., 20th Euromicro International Conference on Parallel, Distributed and Network-Based Processing*, p. 213–220, IEEE, 2012.
- [91] KALMAN, R.; ET AL.: A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering* 82 (1960) 1, p. 35–45.
- [92] KANKAANPÄÄ, P.; PAAVOLAINEN, L.; TIITTA, S.; KARJALAINEN, M.; PÄIVÄRINNE, J.; NIEMINEN, J.; MARJOMÄKI, V.; HEINO, J.; WHITE, D.: BioImageXD: An Open, General-Purpose and High-Throughput Image-Processing Platform. *Nature Methods* 9 (2012) 7, p. 683–689.
- [93] KAUSLER, B. X.; SCHIEGG, M.; ANDRES, B.; LINDNER, M.; LEITTE, H.; HUFNAGEL, L.; KOETHE, U.; HAMPRECHT, F. A.: A Discrete Chain Graph Model for 3D+t Cell Tracking with High Mis-detection Robustness. In: *Proceedings of the European Conference on Computer Vision*, p. 144–157, Florence, Italy: Springer, 2012.

-
- [94] KAYNIG, V.; VAZQUEZ-REINA, A.; KNOWLES-BARLEY, S.; ROBERTS, M.; JONES, T. R.; KASTHURI, N.; MILLER, E.; LICHTMAN, J.; PFISTER, H.: Large-Scale Automatic Reconstruction of Neuronal Processes from Electron Microscopy Images. *Medical image analysis* 22 (2015) 1, p. 77–88.
- [95] KELEMEN, A.; SZÉKELY, G.; GERIG, G.: Elastic Model-based Segmentation of 3D Neuroradiological Data Sets. *IEEE Transactions on Medical Imaging* 18 (1999) 10, p. 828–839.
- [96] KELLER, P.; SCHMIDT, A.; WITTBRODT, J.; STELZER, E.: Reconstruction of Zebrafish Early Embryonic Development by Scanned Light Sheet Microscopy. *Science* 322 (2008) 5904, p. 1065–1069.
- [97] KELLER, P.; SCHMIDT, A.; WITTBRODT, J.; STELZER, E.: Digital Scanned Laser Light-Sheet Fluorescence Microscopy (DSLM) of Zebrafish and Drosophila Embryonic Development. *Cold Spring Harbor Protocols* 2011 (2011) 10, p. 1235–43.
- [98] KELLER, P.; STELZER, E.: Quantitative in vivo Imaging of Entire Embryos with Digital Scanned Laser Light Sheet Fluorescence Microscopy. *Current Opinion in Neurobiology* 18 (2008) 6, p. 624–632.
- [99] KELLER, P. J.: *In Vivo* Imaging of Zebrafish Embryogenesis. *Methods* 62 (2013) 3, p. 268–278.
- [100] KHAN, A.; MIKUT, R.; SCHWEITZER, B.; WEISS, C.; REISCHL, M.: Automatic Tuning of Image Segmentation Parameters by Means of Fuzzy Feature Evaluation. *Advances in Intelligent Systems and Computing* 190 (2013) 5, p. 459–467.
- [101] KHAN, A.; REISCHL, M.; SCHWEITZER, B.; WEISS, C.; MIKUT, R.: Feedback-Driven Design of Normalization Techniques for Biological Images Using Fuzzy Formulation of a Priori Knowledge. *Studies in Computational Intelligence* 445 (2013), p. 167–178.
- [102] KHAN, A.; WEISS, C.; SCHWEITZER, B.; HANSJOSTEN, I.; MIKUT, R.; REISCHL, M.: Multimodal Image Segmentation of Cellular Fragmentation Using Edge Detector and Morphological Operators. *Biomedizinische Technik - Biomedical Engineering* 59 (2014) S1, p. S518–S521.

- [103] KHAN, M. M.; STRACK, S.; WILD, F.; HANASHIMA, A.; GASCH, A.; BROHM, K.; REISCHL, M.; CARNIO, S.; LABEIT, D.; SANDRI, M.; ET AL.: Role of Autophagy, SQSTM1, SH3GLB1, and TRIM63 in the Turnover of Nicotinic Acetylcholine Receptors. *Autophagy* 10 (2014) 1, p. 123–136.
- [104] KHAN, Z.; WANG, Y.-C.; WIESCHAUS, E. F.; KASCHUBE, M.: Quantitative 4D Analyses of Epithelial Folding during *Drosophila* Gastrulation. *Development* 141 (2014) 14, p. 2895–2900.
- [105] KIENDL, H.: *Fuzzy Control methodenorientiert*. München: Oldenbourg, 1997.
- [106] KOBITSKI, A.; OTTE, J. C.; TAKAMIYA, M.; SCHÄFER, B.; MERTES, J.; STEGMAIER, J.; RASTEGAR, S.; RINDONE, F.; HARTMANN, V.; STOTZKA, R.; GARCÍA, A.; WEZEL, J. v.; MIKUT, R.; STRÄHLE, U.; NIENHAUS, G. U.: An Ensemble-Averaged, Cell Density-based Digital Model of Zebrafish Embryo Development Derived from Light-Sheet Microscopy Data with Single-Cell Resolution. *Scientific Reports* 5 (2015) 8601, p. 1–10.
- [107] KOKEL, D.; DUNN, T.; AHRENS, M.; ALSHUT, R.; CHEUNG, C. Y.; SAINT-AMANT, L.; BRUNI, G.; MATEUS, R.; VAN HAM, T.; SHIRAKI, T.; FUKADA, Y.; KOJIMA, D.; YEH, J.-R.; MIKUT, R.; VON LINTIG, J.; ENGERT, F.; ; PETERSON, R.: Identification of Non-visual Photomotor Response Cells in the Vertebrate Hindbrain. *The Journal of Neuroscience* 33 (2013) 9, p. 3834–3843.
- [108] KORZYNSKA, A.; ROSZKOWIAK, L.; LOPEZ, C.; BOSCH, R.; WITKOWSKI, L.; LEJEUNE, M.; ET AL.: Validation of Various Adaptive Threshold Methods of Segmentation Applied to Follicular Lymphoma Digital Images Stained with 3, 3'-Diaminobenzidine&Haematoxylin. *Diagnostic Pathology* 8 (2013) 1, p. 48.
- [109] KRIZHEVSKY, A.; SUTSKEVER, I.; HINTON, G. E.: Imagenet Classification with Deep Convolutional Neural Networks. In: *Advances in Neural Information Processing Systems*, p. 1097–1105, 2012.
- [110] KROLL, A.: *Computational Intelligence*. De Gruyter Oldenbourg, 2016.

-
- [111] KRUSKAL, J. B.: On the Shortest Spanning Subtree of a Graph and the Traveling Salesman Problem. *Proceedings of the American Mathematical Society* 7 (1956) 1, p. 48–50.
- [112] KRZIC, U.; GUNTHER, S.; SAUNDERS, T.; STREICHAN, S.; HUFNAGEL, L.: Multiview Light-Sheet Microscope for Rapid *in toto* Imaging. *Nature Methods* 9 (2012) 7, p. 730–733.
- [113] KVILEKVAL, K.; FEDOROV, D.; OBARA, B.; SINGH, A.; MANJUNATH, B.: Bisque: A Platform for Bioimage Analysis and Management. *Bioinformatics* 26 (2010) 4, p. 544–552.
- [114] KYBIC, J.: Bootstrap Resampling for Image Registration Uncertainty Estimation without Ground Truth. *IEEE Transactions on Image Processing* 19 (2010) 1, p. 64–73.
- [115] KYBIC, J.; UNSER, M.: Fast Parametric Elastic Image Registration. *IEEE Transactions on Image Processing* 12 (2003) 11, p. 1427–1442.
- [116] LE, P. P.; BAHL, A.; UNGAR, L. H.: Using Prior Knowledge to Improve Genetic Network Reconstruction from Microarray Data. *In Silico Biology* 4 (2004) 3, p. 335–353.
- [117] LECUIT, T.; LE GOFF, L.: Orchestrating Size and Shape during Morphogenesis. *Nature* 450 (2007) 7167, p. 189–192.
- [118] LEHMUSSOLA, A.; RUUSUVUORI, P.; SELINUMMI, J.; HUTTUNEN, H.; YLI-HARJA, O.: Computational Framework for Simulating Fluorescence Microscope Images with Cell Populations. *IEEE Transactions on Medical Imaging* 26 (2007) 7, p. 1010–1016.
- [119] LEHMUSSOLA, A.; RUUSUVUORI, P.; SELINUMMI, J.; RAJALA, T.; YLI-HARJA, O.: Synthetic Images of High-Throughput Microscopy for Validation of Image Analysis Methods. *Proceedings of the IEEE* 96 (2008) 8, p. 1348–1360.
- [120] LEHMUSSOLA, A.; SELINUMMI, J.; RUUSUVUORI, P.; NIEMISTO, A.; YLI-HARJA, O.: Simulating Fluorescent Microscope Images of Cell Populations. In: *27th Annual International Conference of the Engineering in Medicine and Biology Society*, p. 3153–3156, IEEE, 2006.

- [121] LEMPITSKY, V.; ZISSERMAN, A.: Learning to Count Objects in Images. In: *Advances in Neural Information Processing Systems*, p. 1324–1332, 2010.
- [122] LEVENTON, M.; GRIMSON, W.; FAUGERAS, O.: Statistical Shape Influence in Geodesic Active Contours. In: *IEEE Conference on Computer Vision and Pattern Recognition, 2000. Proceedings.*, Vol. 1, p. 316–323, IEEE, 2000.
- [123] LI, C.; XU, C.; GUI, C.; FOX, M.: Distance Regularized Level Set Evolution and its Application to Image Segmentation. *IEEE Transactions on Image Processing* 19 (2010) 12, p. 3243–3254.
- [124] LI, G.; LIU, T.; NIE, J.; GUO, L.; MALICKI, J.; MARA, A.; HOLLEY, S. A.; XIA, W.; WONG, S. T.: Detection of Blob Objects in Microscopic Zebrafish Images based on Gradient Vector Diffusion. *Cytometry* 71 (2007) 10, p. 835–845.
- [125] LI, G.; LIU, T.; TAROKH, A.; NIE, J.; GUO, L.; MARA, A.; HOLLEY, S.; WONG, S.: 3D Cell Nuclei Segmentation Based on Gradient Flow Tracking. *BMC Cell Biology* 8 (2007) 1, p. 40.
- [126] LI, H.; MANJUNATH, B.; MITRA, S. K.: Multisensor Image Fusion using the Wavelet Transform. *Graphical Models and Image Processing* 57 (1995) 3, p. 235–245.
- [127] LI, K.; MILLER, E.; CHEN, M.; KANADE, T.; WEISS, L.; CAMPBELL, P.: Cell Population Tracking and Lineage Construction with Spatiotemporal Context. *Medical Image Analysis* 12 (2008) 5, p. 546–566.
- [128] LINDBERG, T.: Feature Detection with Automatic Scale Selection. *International Journal of Computer Vision* 30 (1998) 2, p. 79–116.
- [129] LINKERT, M.; RUEDEN, C.; ALLAN, C.; BUREL, J.; MOORE, W.; PATTERSON, A.; LORANGER, B.; MOORE, J.; NEVES, C.; MACDONALD, D.; TARKOWSKA, A.; STICCO, C.; HILL, E.; ROSSNER, M.; ELICEIRI, K. W.; SWEDLOW, J.: Metadata Matters: Access to Image Data in the Real World. *The Journal of Cell Biology* 189 (2010) 5, p. 777–782.
- [130] LIU, K.; SKIBBE, H.; SCHMIDT, T.; BLEIN, T.; PALME, K.; BROX, T.; RONNEBERGER, O.: Rotation-Invariant HOG Descriptors using Fourier Analysis in Polar and Spherical Coordinates. *International Journal of Computer Vision* 106 (2014) 3, p. 342–364.

-
- [131] LIU, T.; JONES, C.; SEYEDHOSSEINI, M.; TASHIZEN, T.: A Modular Hierarchical Approach to 3D Electron Microscopy Image Segmentation. *Journal of Neuroscience Methods* 226 (2014), p. 88–102.
- [132] LIU, T.; LU, J.; WANG, Y.; CAMPBELL, W.; HUANG, L.; ZHU, J.; XIA, W.; WONG, S.: Computerized Image Analysis for Quantitative Neuronal Phenotyping in Zebrafish. *Journal of Neuroscience Methods* 153 (2006), p. 190–202.
- [133] LIU, T.; NIE, J.; LI, G.; GUO, L.; WONG, S. T. C.: ZFIQ: A Software Package for Zebrafish Biology. *Bioinformatics* 24 (2008) 3, p. 438–439.
- [134] LOU, X.: *Biomedical Data Analysis with Prior Knowledge: Modeling and Learning*. PhD Thesis, University of Heidelberg, 2011.
- [135] LOU, X.; HAMPRECHT, F.: Structured Learning for Cell Tracking. In: *Twenty-Fifth Annual Conference on Neural Information Processing Systems*, p. 1296–1304, 2011.
- [136] LOU, X.; KASTER, F.; LINDNER, M.; KAUSLER, B.; KOTHE, U.; HOCKENDORF, B.; WITTBRODT, J.; JANICKE, H.; HAMPRECHT, F.: Deltr: Digital Embryo Lineage Tree Reconstructor. In: *Proc., IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, p. 1557–1560, IEEE, 2011.
- [137] LOWE, D. G.: Distinctive Image Features from Scale-Invariant Keypoints. *International Journal of Computer Vision* 60 (2004) 2, p. 91–110.
- [138] LUENGO-OROZ, M.; PASTOR-ESCUREDO, D.; CASTRO-GONZALEZ, C.; FAURE, E.; SAVY, T.; LOMBARDOT, B.; RUBIO-GUIVERNAU, J.; DULOQUIN, L.; LEDESMA-CARBAYO, M.; BOURGINE, P.; ET AL.: 3D+t Morphological Processing: Applications to Embryogenesis Image Analysis. *IEEE Transactions on Image Processing* 21 (2012) 8, p. 3518–3530.
- [139] MACKLIN, P.; EDGERTON, M. E.; THOMPSON, A. M.; CRISTINI, V.: Patient-Calibrated Agent-Based Modelling of Ductal Carcinoma In Situ (DCIS): From Microscopic Measurements to Macroscopic Predictions of Clinical Progression. *Journal of Theoretical Biology* 301 (2012), p. 122–140.

- [140] MAES, F.; COLLIGNON, A.; VANDERMEULEN, D.; MARCHAL, G.; SUETENS, P.: Multimodality Image Registration by Maximization of Mutual Information. *IEEE Transactions on Medical Imaging* 16 (1997) 2, p. 187–198.
- [141] MAJI, P.; PAL, S. K.: *Rough-Fuzzy Pattern Recognition: Applications in Bioinformatics and Medical Imaging*, Vol. 3. Hoboken, New Jersey: John Wiley & Sons, 2011.
- [142] MARR, D.; HILDRETH, E.: Theory of Edge Detection. *Proceedings of the Royal Society of London. Series B. Biological Sciences* 207 (1980) 1167, p. 187–217.
- [143] MÄRZ, M.; SCHMIDT, R.; RASTEGAR, S.; STRÄHLE, U.: Regenerative Response Following Stab Injury in the Adult Zebrafish Telencephalon. *Developmental Dynamics* 240 (2011) 9, p. 2221–2231.
- [144] MAŠKA, M.; ULMAN, V.; SVOBODA, D.; MATULA, P.; MATULA, P.; EDERRA, C.; URBIOLA, A.; ESPAÑA, T.; VENKATESAN, S.; BALAK, D. M.; ET AL.: A Benchmark for Comparison of Cell Tracking Algorithms. *Bioinformatics* 30 (2014) 11, p. 1609–1617.
- [145] MAURER JR, C.; QI, R.; RAGHAVAN, V.: A Linear Time Algorithm for Computing Exact Euclidean Distance Transforms of Binary Images in Arbitrary Dimensions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25 (2003) 2, p. 265–270.
- [146] MAZOR, E.; AVERBUCH, A.; BAR-SHALOM, Y.; DAYAN, J.: Interacting Multiple Model Methods in Target Tracking: A Survey. *IEEE Transactions on Aerospace and Electronic Systems* 34 (1998) 1, p. 103–123.
- [147] MCMAHON, A.; REEVES, G. T.; SUPATTO, W.; STATHOPOULOS, A.: Mesoderm Migration in *Drosophila* is a Multi-Step Process Requiring FGF Signaling and Integrin Activity. *Development* 137 (2010) 13, p. 2167–2175.
- [148] MEIJERING, E.; DZYUBACHYK, O.; SMAL, I.; VAN CAPPELLEN, W.: Tracking in Cell and Developmental Biology. In: *Seminars in Cell & Developmental Biology*, Vol. 20, p. 894–902, Elsevier, 2009.

-
- [149] MELANI, C.; PEYRIERAS, N.; MIKULA, K.; ZANELLA, C.; CAMPANA, M.; RIZZI, B.; VERONESI, F.; SARTI, A.; LOMBARDOT, B.; BOURGINE, P.: Cells Tracking in a Live Zebrafish Embryo. In: *IEEE Engineering in Medicine and Biology Magazine*, Vol. 1, p. 1631–1634, 2007.
- [150] MENCATTINI, A.; RABOTTINO, G.; SALICONE, S.; SALMERI, M.: Uncertainty Modeling and Propagation Through RFVs for the Assessment of CADx Systems in Digital Mammography. *IEEE Transactions on Instrumentation and Measurement* 59 (2010) 1, p. 27–38.
- [151] MICHELIN, G.; GUIGNARD, L.; FIUZA, U.-M.; MALANDAIN, G.; ET AL.: Embryo Cell Membranes Reconstruction by Tensor Voting. In: *Proc., IEEE International Symposium on Biomedical Imaging*, p. 1259–1262, 2014.
- [152] MIDGLEY, P.; WEYLAND, M.: 3D Electron Microscopy in the Physical Sciences: The Development of Z-Contrast and EFTEM Tomography. *Ultramicroscopy* 96 (2003) 3, p. 413–431.
- [153] MIERSWA, I.; WURST, M.; KLINKENBERG, R.; SCHOLZ, M.; EULER, T.: YALE: Rapid Prototyping for Complex Data Mining Tasks. In: *Proc., 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (UNGAR, L.; CRAVEN, M.; GUNOPULOS, D.; ELIASSI-RAD, T., Ed.), p. 935–940, New York, NY, USA: ACM, 2006.
- [154] MIKULA, K.; PEYRIÉERAS, N.; REMEŠÍKOVÁ, M.; SMÍŠEK, M.: 4D Numerical Schemes for Cell Image Segmentation and Tracking. *Finite Volumes for Complex Applications VI-Problems & Perspectives* (2011), p. 693–701.
- [155] MIKULA, K.; REMEŠÍKOVÁ, M.; STAŠOVÁ, O.; PEYRIÉERAS, N.: Segmentation and Analysis of 3D Zebrafish Cell Image Data. In: *3rd International Congress on Image and Signal Processing*, Vol. 3, p. 1444–1448, IEEE, 2010.
- [156] MIKUT, R.: *Data Mining in der Medizin und Medizintechnik*. Habilitationsschrift, Universitätsverlag Karlsruhe, 2008.
- [157] MIKUT, R.; BURMEISTER, O.; BRAUN, S.; REISCHL, M.: The Open Source Matlab Toolbox Gait-CAD and its Application to Bioelectric Signal Processing. In: *Proc., DGBMT-Workshop Biosignalverarbeitung, Potsdam*, p. 109–111, 2008.

- [158] MIKUT, R.; DICKMEIS, T.; DRIEVER, W.; GEURTS, P.; HAMPRECHT, F.; KAUSLER, B. X.; LEDESMA-CARBAYO, M. J.; MARÉE, R.; MIKULA, K.; PANTAZIS, P.; RONNEBERGER, O.; SANTOS, A.; STOTZKA, R.; STRÄHLE, U.; PEYRIÉRAS, N.: Automated Processing of Zebrafish Imaging Data - A Survey. *Zebrafish* 10 (2013) 3, p. 401–421.
- [159] MISTELI, T.; SPECTOR, D. L.: Applications of the Green Fluorescent Protein in Cell Biology and Biotechnology. *Nature Biotechnology* 15 (1997) 10, p. 961–964.
- [160] MOOSMANN, J.; ERSHOV, A.; ALTAPOVA, V.; BAUMBACH, T.; PRASAD, M. S.; LABONNE, C.; XIAO, X.; KASHEF, J.; HOFMANN, R.: X-ray Phase-Contrast In Vivo Microtomography Probes new Aspects of Xenopus Gastrulation. *Nature* 497 (2013) 7449, p. 374–377.
- [161] MORRIS, M. K.; SAEZ-RODRIGUEZ, J.; SORGER, P. K.; LAUFFENBURGER, D. A.: Logic-based Models for the Analysis of Cell Signaling Networks. *Biochemistry* 49 (2010) 15, p. 3216–3224.
- [162] MOSALIGANTI, K.; COOPER, L.; SHARP, R.; MACHIRAJU, R.; LEONE, G.; HUANG, K.; SALTZ, J.: Reconstruction of Cellular Biological Structures from Optical Microscopy Data. *IEEE Transactions on Visualization and Computer Graphics* 14 (2008) 4, p. 863–876.
- [163] MOSALIGANTI, K. R.; GELAS, A.; MEGASON, S. G.: An Efficient, Scalable, and Adaptable Framework for Solving Generic Systems of Level-Set PDEs. *Frontiers in Neuroinformatics* 7 (2013), p. 1–14.
- [164] MOSALIGANTI, K. R.; NOCHE, R. R.; XIONG, F.; SWINBURNE, I. A.; MEGASON, S. G.: ACME: Automated Cell Morphology Extractor for Comprehensive Reconstruction of Cell Membranes. *PLoS Computational Biology* 8 (2012) 12, p. e1002780.
- [165] MURPHY, D. B.; DAVIDSON, M. W.: *Fundamentals of Light Microscopy and Electronic Imaging*. Hoboken, New Jersey: Wiley-Blackwell, 2012.
- [166] NAVLAKHA, S.; AHAMMAD, P.; MYERS, E. W.: Unsupervised Segmentation of Noisy Electron Microscopy Images using Salient Watersheds and Region Merging. *BMC Bioinformatics* 14 (2013) 1, p. 294.

-
- [167] OATES, A. C.; GORFINKIEL, N.; GONZÁLEZ-GAITÁN, M.; HEISENBERG, C.-P.: Quantitative Approaches in Developmental Biology. *Nature Reviews Genetics* 10 (2009) 8, p. 517–530.
- [168] OLIVO-MARIN, J.: Extraction of Spots in Biological Images using Multi-scale Products. *Pattern Recognition* 35 (2002) 9, p. 1989–1996.
- [169] OTSU, N.: A Threshold Selection Method from Gray-Level Histograms. *IEEE Transactions on Systems, Man and Cybernetics* 9 (1979), p. 62–66.
- [170] OTTE, J.: *realTox: Real-time Imaging of Toxicant Impact in Whole Organism at Single Cell Resolution*. PhD Thesis, Karlsruhe Institute of Technology, 2013.
- [171] OTTO, F.: DAPI Staining of Fixed Cells for High-Resolution Flow Cytometry of Nuclear DNA. *Methods in Cell Biology* 33 (1990), p. 105–110.
- [172] PAL, N. R.; BHANDARI, D.: On Object Background Classification. *International Journal of Systems Science* 23 (1992) 11, p. 1903–1920.
- [173] PAL, S. K.: Fuzzy Image Processing and Recognition: Uncertainty Handling and Applications. *International Journal of Image and Graphics* 1 (2001) 02, p. 169–195.
- [174] PENG, H.; BRIA, A.; ZHOU, Z.; IANNELLO, G.; LONG, F.: Extensible Visualization and Analysis for Multidimensional Images using Vaa3D. *Nature Protocols* 9 (2014) 1, p. 193–208.
- [175] PENNEY, G. P.; WEESE, J.; LITTLE, J. A.; DESMEDT, P.; HILL, D. L.; ET AL.: A Comparison of Similarity Measures for Use in 2D-3D Medical Image Registration. *IEEE Transactions on Medical Imaging* 17 (1998) 4, p. 586–595.
- [176] PERAVALI, R.; GEHRIG, J.; GISELBRECHT, S.; LÜTJOHANN, D.; HADZHIEV, Y.; MÜLLER, F.; LIEBEL, U.: Automated Feature Detection and Imaging for High-Resolution Screening of Zebrafish Embryos. *BioTechniques* 50 (2011) 5, p. 319–324.
- [177] PERONA, P.; MALIK, J.: Scale-Space and Edge Detection Using Anisotropic Diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12 (1990) 7, p. 629–639.

- [178] PHILOMIN, V.; DURAISWAMI, R.; DAVIS, L.: Pedestrian Tracking from a Moving Vehicle. In: *IEEE Proceedings of the Intelligent Vehicles Symposium*, p. 350–355, IEEE, 2000.
- [179] PINIDIYAARACHCHI, A.; WÄHLBY, C.: Seeded Watersheds for Combined Segmentation and Tracking of Cells. In: *Image Analysis and Processing–ICIAP 2005*, p. 336–343, Springer, 2005.
- [180] PITRONE, P. G.; SCHINDELIN, J.; STUYVENBERG, L.; PREIBISCH, S.; WEBER, M.; ELICEIRI, K. W.; HUISKEN, J.; TOMANCAK, P.: OpenSPIM: An Open-Access Light-Sheet Microscopy Platform. *Nature Methods* 10 (2013), p. 598–599.
- [181] POHORILLE, A.: Incorporating Biological Knowledge into Evaluation of Causal Regulatory Hypotheses. In: *Pacific Symposium on Biocomputing 2003: Kauai, Hawaii, 3-7 January 2003*, p. 128, World Scientific, 2002.
- [182] PREIBISCH, S.; AMAT, F.; STAMATAKI, E.; SAROV, M.; SINGER, R. H.; MYERS, E.; TOMANCAK, P.: Efficient Bayesian-based Multiview Deconvolution. *Nature Methods* 11 (2014) 6, p. 645–648.
- [183] PREIBISCH, S.; ROHLFING, T.; HASAK, M. P.; TOMANCAK, P.: Mosaicing of Single Plane Illumination Microscopy Images using Groupwise Registration and Fast Content-based Image Fusion. *Medical Imaging 2008: Image Processing* 6914 (2008) 1, p. 69140E.
- [184] PREIBISCH, S.; SAALFELD, S.; ROHLFING, T.; TOMANCAK, P.: Bead-based Mosaicing of Single Plane Illumination Microscopy Images using Geometric Local Descriptor Matching. In: *SPIE Medical Imaging*, p. 72592S–72592S, International Society for Optics and Photonics, 2009.
- [185] PREIBISCH, S.; SAALFELD, S.; SCHINDELIN, J.; TOMANCAK, P.: Software for Bead-Based Registration of Selective Plane Illumination Microscopy Data. *Nature Methods* 7 (2010) 6, p. 418–419.
- [186] RADOJEVIC, M.; SMAL, I.; MEIJERING, E.: Automated Neuron Morphology Reconstruction using Fuzzy-logic Detection and Bayesian Tracing Algorithms. In: *Proc., IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, p. 885–888, 2015.

-
- [187] RADOJEVIC, M.; SMAL, I.; NIESSEN, W.; MEIJERING, E.: Fuzzy Logic-based Detection of Neuron Bifurcations in Microscopy Images. In: *Proc., IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, p. 1307–1310, 2014.
- [188] RAND, W. M.: Objective Criteria for the Evaluation of Clustering Methods. *Journal of the American Statistical Association* 66 (1971) 336, p. 846–850.
- [189] RAPOPORT, D. H.; BECKER, T.; MAMLOUK, A. M.; SCHICKTANZ, S.; KRUSE, C.: A Novel Validation Algorithm Allows for Automated Cell Tracking and the Extraction of Biologically Meaningful Parameters. *PLoS ONE* 6 (2011) 11, p. e27315.
- [190] REISCHL, M.; ALSHUT, R.; MIKUT, R.: On Robust Feature Extraction and Classification of Inhomogeneous Data Sets. In: *Proc., 20. Workshop Computational Intelligence*, p. 124–143, KIT Scientific Publishing, 2010.
- [191] RUBIO-GUIVERNAU, J.; GURCHENKOV, V.; LUENGO-OROZ, M.; DU-LOQUIN, L.; BOURGINE, P.; SANTOS, A.; PEYRIERAS, N.; LEDESMA-CARBAYO, M.: Wavelet-Based Image Fusion in Multi-View Three-Dimensional Microscopy. *Bioinformatics* 28 (2012) 2, p. 238–245.
- [192] RUDZKI, M.: Vessel Detection Method Based on Eigenvalues of the Hessian Matrix and its Applicability to Airway Tree Segmentation. In: *Proceedings of the 11th International PhD Workshop*, p. 100–105, 2009.
- [193] RUEDEN, C.; ELICEIRI, K. W.; WHITE, J. G.: VisBio: A Computational Tool for Visualization of Multidimensional Biological Image Data. *Traffic* 5 (2004) 6, p. 411–417.
- [194] RUSU, R. B.; COUSINS, S.: 3D is Here: Point Cloud Library (PCL). In: *IEEE International Conference on Robotics and Automation*, p. 1–4, Shanghai, China, 2011.
- [195] RUUSUVUORI, P.; LEHMUSSOLA, A.; SELINUMMI, J.; RAJALA, T.; HUTTUNEN, H.; YLI-HARJA, O.: Benchmark Set of Synthetic Images for Validating Cell Image Analysis Algorithms. In: *Proceedings of the 16th European Signal Processing Conference*, 2008.

- [196] SANTO, M.; LIGUORI, C.; PAOLILLO, A.; PIETROSANTO, A.: Standard Uncertainty Evaluation in Image-Based Measurements. *Measurement* 36 (2004) 3, p. 347–358.
- [197] SBALZARINI, I. F.; KOUMOUTSAKOS, P.: Feature Point Tracking and Trajectory Analysis for Video Imaging in Cell Biology. *Journal of Structural Biology* 151 (2005) 2, p. 182–195.
- [198] SCHIEGG, M.; HANSLOVSKY, P.; HAUBOLD, C.; KOETHE, U.; HUFNAGEL, L.; HAMPRECHT, F. A.: Graphical Model for Joint Segmentation and Tracking of Multiple Dividing Cells. *Bioinformatics* (2014), p. btu764.
- [199] SCHIEGG, M.; HANSLOVSKY, P.; KAUSLER, B. X.; HUFNAGEL, L.; HAMPRECHT, F. A.: Conservation Tracking. In: *IEEE International Conference on Computer Vision*, p. 2928–2935, IEEE, 2013.
- [200] SCHINDELIN, J.; ARGANDA-CARRERAS, I.; FRISE, E.; KAYNIG, V.; LONGAIR, M.; PIETZSCH, T.; PREIBISCH, S.; RUEDEN, C.; SAALFELD, S.; SCHMID, B.; TINEVEZ, J.-Y.; WHITE, D. J.; HARTENSTEIN, V.; ELICEIRI, K.; TOMANCAK, P.; CARDONA, A.: Fiji: An Open-Source Platform for Biological-Image Analysis. *Nature Methods* 9 (2012) 7, p. 676–682.
- [201] SCHROEDER, W.; MARTIN, K.; AVILA, L. S.; LAW, C. C.: The Visualization Toolkit User’s Guide, Version 4.0. *Kitware* 4 (2001).
- [202] SE, S.; LOWE, D.; LITTLE, J.: Vision-Based Mobile Robot Localization and Mapping using Scale-Invariant Features. In: *IEEE International Conference on Robotics and Automation*, Vol. 2, p. 2051–2058, IEEE, 2001.
- [203] SEZGIN, M.; SANKUR, B.: Survey over Image Thresholding Techniques and Quantitative Performance Evaluation. *Journal of Electronic Imaging* 13 (1) (2004), p. 146–165.
- [204] SHAHID, M.; TAKAMIYA, M.; STEGMAIER, J.; MIDDEL, V.; GRADL, M.; KLÜVER, N.; MIKUT, R.; DICKMEIS, T.; SCHOLZ, S.; RASTEGAR, S.; YANG, L.; STRÄHLE, U.: Zebrafish Biosensor for Toxicant Induced Muscle Hyperactivity. *Scientific Reports* 6 (2016) 23768, p. 1–14.
- [205] SHARMA, H.; ALEKSEYCHUK, A.; LESKOVSKY, P.; HELLWICH, O.; ANAND, R.; ZERBE, N.; HUFNAGL, P.; ET AL.: Determining Similarity


- in Histological Images using Graph-Theoretic Description and Matching Methods for Content-based Image Retrieval in Medical Diagnostics. *Diagnostic Pathology* 7 (2012) 1, p. 134.
- [206] SHIN, J.; POLING, J.; PARK, H.-C.; APPEL, B.: Notch Signaling Regulates Neural Precursor Allocation and Binary Neuronal Fate Decisions in Zebrafish. *Development* 134 (2007) 10, p. 1911–1920.
- [207] SIANG TAN, K.; MAT ISA, N. A.: Color Image Segmentation using Histogram Thresholding–Fuzzy C-Means Hybrid Approach. *Pattern Recognition* 44 (2011) 1, p. 1–15.
- [208] SINTORN, I. M.; BISCHOF, L.; JACKWAY, P.; HAGGARTY, S.; BUCKLEY, M.: Gradient Based Intensity Normalization. *Journal of Microscopy* 240 (2010) 3, p. 249–258.
- [209] SMITH, S. W.; ET AL.: *The Scientist and Engineer's Guide to Digital Signal Processing*. California Technical Pub. San Diego, 1997.
- [210] SOILLE, P.: *Morphological Image Analysis: Principles and Applications*. Springer-Verlag New York, Inc. Secaucus, NJ, USA, 2003.
- [211] SOMMER, C.; STRAEHLE, C.; KOTHE, U.; HAMPRECHT, F.: ilastik: Interactive Learning and Segmentation Toolkit. In: *Proc., IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, p. 230–233, IEEE, 2011.
- [212] SONG, B.; CHOI, H.; LEE, H. S.: Surveillance Tracking System using Passive Infrared Motion Sensors in Wireless Sensor Network. In: *International Conference on Information Networking*, p. 1–5, IEEE, 2008.
- [213] SORZANO, C. O.; THEVENAZ, P.; UNSER, M.: Elastic Registration of Biological Images using Vector-Spline Regularization. *IEEE Transactions on Biomedical Engineering* 52 (2005) 4, p. 652–663.
- [214] SOUBIÈS, E.; WEISS, P.; DESCOMBES, X.; ET AL.: A 3D Segmentation Algorithm for Ellipsoidal Shapes. Application to Nuclei Extraction. *HAL-Inria* (2012), p. 1–9.

- [215] STEGMAIER, J.; ALSHUT, R.; REISCHL, M.; MIKUT, R.: Information Fusion of Image Analysis, Video Object Tracking, and Data Mining of Biological Images using the Open Source MATLAB Toolbox Gait-CAD. *Biomedizinische Technik (Biomedical Engineering)* 57 (S1) (2012), p. 458–461.
- [216] STEGMAIER, J.; AMAT, F.; LEMON, W. B.; MCDOLE, K.; WAN, Y.; TEODORO, G.; MIKUT, R.; KELLER, P. J.: Real-Time Three-Dimensional Cell Segmentation in Large-Scale Microscopy Data of Developing Embryos. *Developmental Cell* 36 (2016) 2, p. 225–240.
- [217] STEGMAIER, J.; ARZ, J.; SCHOTT, B.; OTTE, J. C.; KOBITSKI, A.; NIENHAUS, G. U.; STRÄHLE, U.; SANDERS, P.; MIKUT, R.: Generating Semi-Synthetic Validation Benchmarks for Embryomics. In: *Proc., IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2016.
- [218] STEGMAIER, J.; KHAN, A.; REISCHL, M.; MIKUT, R.: Challenges of Uncertainty Propagation in Image Analysis. In: *Proc., 22. Workshop Computational Intelligence, Dortmund*, p. 55–69, 2012.
- [219] STEGMAIER, J.; MIKUT, R.: Fuzzy-based Propagation of Prior Knowledge to Improve Large-Scale Image Analysis Pipelines. arxiv.org, submitted paper, 2016.
- [220] STEGMAIER, J.; OTTE, J. C.; KOBITSKI, A.; BARTSCHAT, A.; GARCIA, A.; NIENHAUS, G. U.; STRÄHLE, U.; MIKUT, R.: Fast Segmentation of Stained Nuclei in Terabyte-Scale, Time Resolved 3D Microscopy Image Stacks. *PLoS ONE* 9 (2014) 2, p. e90036.
- [221] STEGMAIER, J.; SHAHID, M.; TAKAMIYA, M.; YANG, L.; RASTEGAR, S.; REISCHL, M.; STRÄHLE, U.; MIKUT, R.: Automated Prior Knowledge-Based Quantification of Neuronal Patterns in the Spinal Cord of Zebrafish. *Bioinformatics* 30 (2014) 5, p. 726–733.
- [222] STRAEHLE, C.; KOETHE, U.; KNOTT, G.; BRIGGMAN, K.; DENK, W.; HAMPRECHT, F. A.: Seeded Watershed Cut Uncertainty Estimators for Guided Interactive Segmentation. In: *IEEE Conference on Computer Vision and Pattern Recognition*, p. 765–772, IEEE, 2012.

-
- [223] STRAW, A. D.; BRANSON, K.; NEUMANN, T. R.; DICKINSON, M. H.: Multi-Camera Real-Time Three-Dimensional Tracking of Multiple Flying Animals. *Journal of The Royal Society Interface* (2010), p. rsif20100230.
- [224] SVOBODA, D.; KAŠÍK, M.; MAŠKA, M.; HUBENÝ, J.; STEJSKAL, S.; ZIMMERMANN, M.: On Simulating 3D Fluorescent Microscope Images. In: *Computer Analysis of Images and Patterns*, p. 309–316, Springer, 2007.
- [225] SVOBODA, D.; KOZUBEK, M.; STEJSKAL, S.: Generation of Digital Phantoms of Cell Nuclei and Simulation of Image Formation in 3D Image Cytometry. *Cytometry Part A* 75 (2009) 6, p. 494–509.
- [226] SVOBODA, D.; ULMAN, V.: Generation of Synthetic Image Datasets for Time-lapse Fluorescence Microscopy. In: *Image Analysis and Recognition*, p. 473–482, Springer, 2012.
- [227] TAKAMIYA, M.; STEGMAIER, J.; KOBITSKI, A.; WEGER, B.; YANG, L.; SORGE, S.; OTTE, J. C.; SCHÄFER, B.; HARTMANN, V.; WEZEL, J. V.; STOTZKA, R.; NIENHAUS, G. U.; MIKUT, R.; DICKMEIS, T.; RASTEGAR, S.; STRÄHLE, U.: Distinct Cranial Neural Crest Populations Contribute to the Development of the Eye. *in revision* (2016).
- [228] TARI, L.; BARAL, C.; KIM, S.: Fuzzy C-Means Clustering with Prior Biological Knowledge. *Journal of Biomedical Informatics* 42 (2009) 1, p. 74–81.
- [229] TIZHOOSH, H. R.: Image Thresholding using Type II Fuzzy Sets. *Pattern Recognition* 38 (2005) 12, p. 2363–2372.
- [230] TOMER, R.; KHAIRY, K.; AMAT, F.; KELLER, P.: Quantitative High-Speed Imaging of Entire Developing Embryos with Simultaneous Multiview Light-Sheet Microscopy. *Nature Methods* (2012), p. 755–763.
- [231] TSIEN, R. Y.: The Green Fluorescent Protein. *Annual Review of Biochemistry* 67 (1998) 1, p. 509–544.
- [232] UMEYAMA, S.: Least-squares Estimation of Transformation Parameters Between Two Point Patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 13 (1991) 4, p. 376–380.
- [233] VACHIER, C.; MEYER, F.: The Viscous Watershed Transform. *Journal of Mathematical Imaging and Vision* 22 (2005) 2, p. 251–267.

- [234] VIOLA, P.; JONES, M.: Rapid Object Detection using a Boosted Cascade of Simple Features. In: *IEEE Conference on Computer Vision and Pattern Recognition*, Vol. 1, p. I-511, IEEE, 2001.
- [235] VOIGT, C.; ADAMY, J.: *Formelsammlung der Matrizenrechnung*. Oldenbourg Verlag, 2007.
- [236] VOIGTMAN, E.: Comparison of Signal-to-Noise Ratios. *Analytical Chemistry* 69 (1997) 2, p. 226–234.
- [237] VOVK, U.; PERNUS, F.; LIKAR, B.: A Review of Methods for Correction of Intensity Inhomogeneity in MRI. *IEEE Transactions on Medical Imaging* 26 (2007) 3, p. 405–421.
- [238] VU, N.; MANJUNATH, B.: Shape Prior Segmentation of Multiple Objects with Graph Cuts. In: *IEEE Conference on Computer Vision and Pattern Recognition*, p. 1–8, IEEE, 2008.
- [239] WANG, W.; HUANG, X.; ESMAILI, A.: Texture-Based Foam Segmentation and Analysis. *Industrial & Engineering Chemistry Research* 50 (2011) 10, p. 6071–6081.
- [240] WANG, X.-Y.; WANG, T.; BU, J.: Color Image Segmentation using Pixel Wise Support Vector Machine Classification. *Pattern Recognition* 44 (2011) 4, p. 777–787.
- [241] WARD JR, J. H.: Hierarchical Grouping to Optimize an Objective Function. *Journal of the American Statistical Association* 58 (1963) 301, p. 236–244.
- [242] WASSENBERG, J.; MIDDELMANN, W.; SANDERS, P.: Highly Efficient Screening for Point-Like Targets via Concentric Shells. In: *Advanced Maui Optical and Space Surveillance Technologies Conference*, Vol. 1, p. 73, 2010.
- [243] WEBER, M.; HUISKEN, J.: Light Sheet Microscopy for Real-Time Developmental Biology. *Current Opinion in Genetics & Development* 21(5) (2011), p. 566–572.
- [244] WEGER, M.; WEGER, B.; TAKAMIYA, M.; STEGMAIER, J.; KOBITSKI, A.; NIENHAUS, G. U.; MIKUT, R.; STRÄHLE, U.; DICKMEIS, T.: Regulation of Cholesterol Biogenesis by the Glucose-sensing Transcription Factor MondoA is Required for Zebrafish Epiboly. *under review* (2014).

-
- [245] WEISBERG, S. P.; MCCANN, D.; DESAI, M.; ROSENBAUM, M.; LEIBEL, R. L.; FERRANTE, A. W.; ET AL.: Obesity is Associated with Macrophage Accumulation in Adipose Tissue. *Journal of Clinical Investigation* 112 (2003) 12, p. 1796–1808.
- [246] WELVAERT, M.; ROSSEEL, Y.: On the Definition of Signal-to-Noise Ratio and Contrast-to-Noise Ratio for fMRI Data. *PLoS ONE* 8 (2013) 11, p. e77089.
- [247] WILLIAMS, C.: Jaccard's Generic Coefficient and Coefficient of Floral Community, in Relation to the Logarithmic Series and the Index of Diversity. *Annals of Botany* 13 (1949) 49, p. 53–58.
- [248] WILSON, R.; GRANLUND, G.: The Uncertainty Principle in Image Processing. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6 (1984), p. 758–767.
- [249] WINDHAGER, L.; ERHARD, F.; ZIMMER, R.: Fuzzy Modeling. In: *Modeling in Systems Biology*, p. 179–205, Springer, 2011.
- [250] WOLF, S.; LOOSE, T.; SCHABLOWSKI, M.; DÖDERLEIN, L.; RUPP, R.; GERNER, H. J.; BRETTHAUER, G.; MIKUT, R.: Automated Feature Assessment in Instrumented Gait Analysis. *Gait & Posture* 23(3) (2006), p. 331–338.
- [251] YU, W.; LEE, H.; HARIHARAN, S.; BU, W.; AHMED, S.: Quantitative Neurite Outgrowth Measurement Based on Image Segmentation with Topological Dependence. *Cytometry Part A* 75 (2009) 4, p. 289–297.
- [252] ZADEH, L.: Fuzzy Sets. *Information and Control* 8 (1965), p. 338–353.
- [253] ZHOU, X.; GUPTA, A.; COMANICIU, D.: An Information Fusion Framework for Robust Shape Tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 27 (2005) 1, p. 115–129.
- [254] ZHUANG, H.; LOW, K.-S.; YAU, W.-Y.: Multichannel Pulse-Coupled-Neural-Network-based Color Image Segmentation for Object Detection. *IEEE Transactions on Industrial Electronics* 59 (2012) 8, p. 3299–3308.



New Methods to Improve Large-Scale Microscopy Image Analysis with Prior Knowledge and Uncertainty

Multidimensional imaging techniques provide powerful ways to examine various kinds of scientific questions. The routinely produced datasets in the terabyte-range, however, can hardly be analyzed manually and require an extensive use of automated image analysis. The present work introduces a new concept for the estimation and propagation of uncertainty involved in image analysis operators and new segmentation algorithms that are suitable for terabyte-scale analyses of 3D+t microscopy images.



Gedruckt auf FSC-zertifiziertem Papier